Configuring IP Services

Router Software Version 11.0
Site Manager Software Version 5.0

Part No. 114065 Rev. A
August 1996

Bay Networks

Bay Networks

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Copyright © 1988–1996 Bay Networks, Inc.

All rights reserved. Printed in the USA. August 1996.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Bay Networks, Inc.

The software described in this document is furnished under a license agreement and may only be used in accordance with the terms of that license. A summary of the Software License is included in this document.

Restricted Rights Legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notice for All Other Executive Agencies

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Trademarks of Bay Networks, Inc.

ACE, AFN, AN, BCN, BLN, BN, BNX, CN, FN, FRE, GAME, LN, Optivity, PPX, SynOptics, SynOptics Communications, Wellfleet and the Wellfleet logo are registered trademarks and ANH, ASN, Bay•SIS, BCNX, BLNX, EZ Install, EZ Internetwork, EZ LAN, PathMan, PhonePlus, Quick2Config, RouterMan, SPEX, Bay Networks, Bay Networks Press, the Bay Networks logo and the SynOptics logo are trademarks of Bay Networks, Inc.

Third-Party Trademarks

All other trademarks and registered trademarks are the property of their respective owners.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Bay Networks, Inc. reserves the right to make changes to the products described in this document without notice.

Bay Networks, Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product are Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Bay Networks Software License

Note: This is Bay Networks basic license document. In the absence of a software license agreement specifying varying terms, this license -- or the license included with the particular product -- shall govern licensee's use of Bay Networks software.

This Software License shall govern the licensing of all software provided to licensee by Bay Networks ("Software"). Bay Networks will provide licensee with Software in machine-readable form and related documentation ("Documentation"). The Software provided under this license is proprietary to Bay Networks and to third parties from whom Bay Networks has acquired license rights. Bay Networks will not grant any Software license whatsoever, either explicitly or implicitly, except by acceptance of an order for either Software or for a Bay Networks product ("Equipment") that is packaged with Software. Each such license is subject to the following restrictions:

1. Upon delivery of the Software, Bay Networks grants to licensee a personal, nontransferable, nonexclusive license to use the Software with the Equipment with which or for which it was originally acquired, including use at any of licensee's facilities to which the Equipment may be transferred, for the useful life of the Equipment unless earlier terminated by default or cancellation. Use of the Software shall be limited to such Equipment and to such facility. Software which is licensed for use on hardware not offered by Bay Networks is not subject to restricted use on any Equipment, however, unless otherwise specified on the Documentation, each licensed copy of such Software may only be installed on one hardware item at any time.
2. Licensee may use the Software with backup Equipment only if the Equipment with which or for which it was acquired is inoperative.
3. Licensee may make a single copy of the Software (but not firmware) for safekeeping (archives) or backup purposes.
4. Licensee may modify Software (but not firmware), or combine it with other software, subject to the provision that those portions of the resulting software which incorporate Software are subject to the restrictions of this license. Licensee shall not make the resulting software available for use by any third party.
5. Neither title nor ownership to Software passes to licensee.
6. Licensee shall not provide, or otherwise make available, any Software, in whole or in part, in any form, to any third party. Third parties do not include consultants, subcontractors, or agents of licensee who have licensee's permission to use the Software at licensee's facility, and who have agreed in writing to use the Software only in accordance with the restrictions of this license.
7. Third-party owners from whom Bay Networks has acquired license rights to software that is incorporated into Bay Networks products shall have the right to enforce the provisions of this license against licensee.
8. Licensee shall not remove or obscure any copyright, patent, trademark, trade secret, or similar intellectual property or restricted rights notice within or affixed to any Software and shall reproduce and affix such notice on any backup copy of Software or copies of software resulting from modification or combination performed by licensee as permitted by this license.

Bay Networks Software License *(continued)*

9. Licensee shall not reverse assemble, reverse compile, or in any way reverse engineer the Software. [Note: For licensees in the European Community, the Software Directive dated 14 May 1991 (as may be amended from time to time) shall apply for interoperability purposes. Licensee must notify Bay Networks in writing of any such intended examination of the Software and Bay Networks may provide review and assistance.]
10. Notwithstanding any foregoing terms to the contrary, if licensee licenses the Bay Networks product "Site Manager," licensee may duplicate and install the Site Manager product as specified in the Documentation. This right is granted solely as necessary for use of Site Manager on hardware installed with licensee's network.
11. This license will automatically terminate upon improper handling of Software, such as by disclosure, or Bay Networks may terminate this license by written notice to licensee if licensee fails to comply with any of the material provisions of this license and fails to cure such failure within thirty (30) days after the receipt of written notice from Bay Networks. Upon termination of this license, licensee shall discontinue all use of the Software and return the Software and Documentation, including all copies, to Bay Networks.
12. Licensee's obligations under this license shall survive expiration or termination of this license.

Contents

Configuring IP Services

About This Guide

Before You Begin xix
Conventions xx
Acronyms xxi
Ordering Bay Networks Publications xxii

Technical Support and Online Services

Bay Networks Customer Service xxiv
Bay Networks Information Services xxv
World Wide Web xxv
Customer Service FTP xxv
Support Source CD xxvi
CompuServe xxvi
InfoFACTS xxvii
How to Get Help xxvii

Chapter 1
IP Concepts, Terminology, and Features

IP Addresses 1-2
Subnet Addressing 1-4
Supernet Addressing 1-7
Classless Interdomain Routing 1-8
Autonomous Systems 1-8
Routing Protocols 1-8
Routing Information Protocol (RIP) 1-9
Open Shortest Path First (OSPF) Protocol 1-10
Border Gateway Protocol (BGP) 1-10
Exterior Gateway Protocol (EGP) 1-10

Route Preferences 1-11
Route Weights 1-11
IP Routing Policies and Filters 1-13
IP Traffic Filters 1-17
RFC Compliance 1-17

Chapter 2
Starting IP Services

Starting IP 2-3
Deleting IP from an Interface 2-4
Customizing IP 2-4
Starting RIP 2-5
Adding RIP to an IP Interface 2-5
Deleting RIP from An IP Interface 2-6
Customizing RIP 2-6
Starting OSPF 2-7
Adding OSPF to an IP Interface 2-8
Deleting OSPF from an IP Interface 2-8
Customizing OSPF 2-8
Starting BGP 2-9
Deleting BGP from the Router 2-10
Deleting BGP-3 and BGP-4 from the Router 2-11
Customizing BGP 2-11
Starting EGP 2-12
Deleting EGP from the Router 2-12
Customizing EGP 2-13
Starting NetBIOS over IP 2-13
Adding NetBIOS to an IP Interface 2-13
Deleting NetBIOS from an IP Interface 2-14
Customizing NetBIOS over IP 2-14
Using the Circuitless IP Interface 2-15
Starting IP on the Circuitless Interface 2-15
Adding OSPF and BGP to the Circuitless Interface 2-16
Selecting Slots to Support the Circuitless Interface 2-16
Configuring an Unnumbered IP Interface 2-16
Using the Alternate Associated Address Option 2-18

Chapter 3
Customizing IP Services

Configuring IP on the Router 3-3
Configuring the Router for Host-Only Mode 3-4
Setting the Time-to-Live Value on a Source Packet 3-8
Setting the RIP Diameter 3-8
Allowing an All-Zero or All-One Subnet Address 3-9
Estimating the Size of the Routing Table 3-9
Using a Default Route for an Unknown Subnet 3-10
Specifying the Maximum Number of IP Policies 3-10
Disabling Route Filter Support 3-10
Enabling Equal-Cost Multipath Support 3-11
Enabling ISP Mode on the Router 3-12
Configuring an IP Interface 3-13
Configuring a Multinet Interface 3-15
Specifying a Broadcast Address for an Interface 3-16
Specifying a Subnet Broadcast Address 3-16
Specifying the Cost of an Interface 3-17
Enabling MTU Discovery on an Interface 3-18
Configuring ICMP Features on an Interface 3-19
Enabling ICMP Address-Mask Replies on an Interface 3-19
Enabling ICMP Redirect Messages 3-19
Enabling All-Subnet Broadcasting on an Interface 3-20
Disabling UDP Checksum Processing on the Interface 3-21
Specifying a MAC Address or E.164 Address 3-21
Enabling Source Routing over a Token Ring Network 3-22
Configuring an SMDS Address 3-24
Configuring a WAN Address for a Frame Relay Network 3-24
Specifying the Maximum Size of the Forwarding Table 3-25
Configuring an Interface for an ATM Logical IP Subnet 3-26
Configuring Address Resolution 3-27
Selecting an Address Resolution Scheme for an Interface 3-29
Selecting an Encapsulation Option for ARP and Probe 3-30
Enabling Proxy ARP on an Interface 3-31
Timing Out Entries in the Address Resolution Cache 3-32

Defining an Adjacent Host 3-33
Configuring IP Accounting on a Frame Relay Interface 3-35
Enabling IP Accounting on the Router 3-35
Specifying the Maximum Size of the IP Accounting Table 3-35
Controlling Notification of a Full IP Accounting Table 3-36
Copying the IP Accounting Table to the Checkpoint Table 3-36
Configuring the Trivial File Transfer Protocol 3-36
Configuring the Revised IP Security Option on an Interface 3-37
Security Label Format 3-39
Inbound IP Datagrams 3-41
Forwarded IP Datagrams 3-42
Originated IP Datagrams 3-42
Unlabeled IP Datagrams 3-42
RIPSO Example 3-43
Defining a Static Route 3-45
Defining a Black Hole for a Supernet 3-46
Configuring Router Discovery 3-46
Connecting the Router to a Blacker Front End 3-47
BFE Addressing 3-48
Configuring Blacker Front End Support 3-49

Chapter 4
Customizing RIP Services

Enabling and Disabling RIP on an Interface 4-2
Selecting the RIP Version 4-2
Sending RIP Updates on an Interface 4-3
Sending Triggered Updates 4-4
Specifying a Time-to-Live Value 4-5
Receiving RIP Updates on an Interface 4-5
Authenticating the Password on a Version 2 Update 4-6
Supplying a Default Route on an Interface 4-7
Listening for a Default Route 4-7
Configuring RIP for Equal-Cost Multipath Support 4-8
Configuring a RIP Interface for Dial-Optimized Routing 4-9
Setting RIP Timers on an Interface 4-9
Specifying an Update Interval 4-9

Specifying a Timeout Period 4-10
Specifying a Holddown Period 4-10
Specifying a Stabilization Time 4-11

Chapter 5
Customizing OSPF Services

OSPF Concepts and Terminology 5-3
OSPF Addresses and Variable-Length Masks 5-3
OSPF Neighbors 5-3
Neighbor Adjacencies 5-4
Designated Routers 5-4
OSPF Areas 5-5
Area Internal Routers 5-6
Area Border Router 5-6
Area Identifiers 5-6
Boundary Routers and AS External Routes 5-6
OSPF Implementation Notes 5-7
Configuring OSPF Global Features 5-9
Enabling and Disabling OSPF on the Router 5-9
Supplying an OSPF ID 5-9
Configuring the Soloist and Backup Soloist on a Slot 5-10
Using the Holddown Timer to Relieve the CPU 5-11
Configuring Message Logging 5-11
Configuring OSPF for Equal-Cost Multipath Routing 5-13
Configuring OSPF on an IP Interface 5-14
Enabling and Disabling OSPF 5-14
Configuring an Area ID 5-14
Specifying the Network Type 5-15
Specifying Router Priority for a Multiaccess Network 5-17
Estimating the Transit Delay 5-18
Configuring Interface Timers 5-19
Setting the Retransmit Interval 5-19
Setting the Hello Interval 5-20
Setting the Dead Interval 5-21
Poll Interval for NBMA Neighbors 5-22
Specifying Metric Cost 5-23

Specifying the MTU Size 5-24
Configuring OSPF on a Circuitless Interface 5-25
Configuring a Neighbor on an NBMA Interface 5-25
Configuring a Boundary Router 5-25
Enabling the Boundary Function 5-25
Configuring AS External Advertisements 5-26
Using the External Route Tag in an ASE 5-28
Configuring Authentication on an Interface 5-30
Configuring an Area Border Router 5-30
Configuring a Virtual Backbone Link through a Transit Area 5-32
Defining a Range for a Summary Route 5-34
Configuring Border Services for a Stub 5-34
Applying IP Policies to OSPF Interfaces 5-35

Chapter 6
Customizing BGP Services

BGP Concepts and Terminology 6-2
Interior BGP in Intra-AS Routing 6-3
IBGP in a Transit AS 6-3
BGP Path Attributes 6-5
BGP-4 Local Preference Values 6-7
BGP Implementation Notes 6-8
Configuring BGP on the Router 6-9
Enabling and Disabling BGP 6-9
Supplying a BGP Identifier 6-9
Identifying the Local AS 6-10
Configuring BGP for Intra-AS Routing 6-10
Setting the Interval Timer 6-11
Allowing Redundant Connections 6-11
Enabling Multihop Connections 6-12
Disabling Dynamic Policy Configuration 6-12
Configuring BGP as a Soloist 6-13
Establishing a Peer-to-Peer Connection 6-13
Defining the Peer-to-Peer Connection 6-13
Using the Circuitless IP Interface for a Peer-to-Peer Connection 6-14
Configuring BGP Peers over an Unnumbered Point-to-Point Link 6-14

Initiating a Peer-to-Peer Connection 6-16
Negotiating the BGP Version 6-17
Keeping the Connection Alive 6-17
Specifying an Update Interval for the Connection 6-18
Specifying a Holddown Time 6-19
Setting a Minimum AS Origination Interval 6-19
Overriding the Local AS Number 6-20
Specifying a Maximum Update Size 6-20
Setting the Route Echo Switch 6-20
Designing a Route Selection Strategy 6-21
Assigning AS Weight and Class Values 6-21
Configuring Routing Policies 6-22
Configuring a BGP-4 AS Pattern-Matching Policy 6-23
Best Route Calculation for Equal Routes 6-24
OSPF/BGP Interaction 6-25
Configuring BGP Message Logging 6-25

Chapter 7
Customizing EGP Services

EGP Concepts and Terminology 7-1
EGP Implementation Notes 7-5
Configuring EGP on the Router 7-6
Enabling and Disabling EGP 7-6
Supplying a Local AS Number 7-6
Configuring a Neighbor 7-6
Specifying the Neighbor's Address 7-7
Specifying the Gateway Mode 7-7
Enabling and Disabling the Neighbor Relationship 7-8
Selecting the Acquisition Mode 7-8
Selecting the Poll Mode 7-8
Setting Neighbor Timers 7-9

Chapter 8
Customizing NetBIOS over IP

NetBIOS Concepts and Terminology 8-2
NetBIOS in an IP Environment 8-3

Forwarding Name Queries over an Unnumbered Interface 8-6
Enabling and Disabling NetBIOS on a Router 8-6
Specifying a TTL Value for a Rebroadcast Packet 8-6
Enabling the Insertion of Record Route Option 8-7
Configuring a NetBIOS Cache 8-7
Enabling Name Caching on the Router 8-8
Creating a MIB Instance for a Cached Name 8-8
Specifying the Size of the Name Cache 8-8
Aging a Cache Entry 8-9
Customizing a Cache Search 8-10
Configuring NetBIOS on an IP Interface 8-11
Enabling and Disabling NetBIOS on an Interface 8-11
Enabling Name Caching on the Interface 8-11
Disabling Inbound and Outbound Broadcasts 8-11
Supplying a Rebroadcast Address 8-12
Configuring a Static NetBIOS Name and Address 8-12
Enabling and Disabling Static Name Caching 8-12
Specifying the NetBIOS Station Name 8-12
Specifying the NetBIOS Station Address 8-13
Supplying a Scope Identifier 8-13
Adding a Traffic Filter to a NetBIOS Interface 8-14

Appendix A
Site Manager Parameters

BGP Parameters A-1
BGP Configuration Parameters A-1
BGP Global Parameters A-2
BGP-3 Global Parameters A-5
BGP-4 Global Parameters A-6
BGP Peer Parameters A-6
BGP AS Weight and Weight Class Parameters A-12
BGP Event Message Parameters A-16
EGP Parameters A-18
EGP Global Parameters A-18
EGP Neighbor Parameters A-19
IP Parameters A-22

IP Configuration Parameters A-22
IP Interface Parameters A-24
IP Global Parameters A-39
IP Accounting Parameters A-46
Static Route Parameters A-47
Adjacent Host Parameters A-51
TFTP Parameters A-53
RIPSO Parameters A-55
Router Discovery Parameters A-64
NetBIOS A-66
NetBIOS Global Parameters A-66
NetBIOS/IP Interface Table Parameters A-69
Static Entry Table Parameters A-71
OSPF Parameters A-73
OSPF Global Parameters A-73
OSPF Interface Paramters A-78
Neighbor Parameters for an NBMA Interface A-84
OSPF Area Parameters A-85
Area Range Parameters A-88
OSPF Virtual Interface Parameters A-90
RIP Parameters A-94

Appendix B
Routing Policies

RIP-Specific Accept Policy Parameters B-7
OSPF-Specific Accept Policy Parameters B-8
EGP-Specific Accept Policy Parameters B-9
BGP-3-Specific Accept Policy Parameters B-11
BGP-4-Specific Accept Policy Parameters B-15
IP Announce Policy Parameters B-20
RIP-Specific Announce Policy Parameters B-38
OSPF-Specific Announce Policy Parameters B-39
EGP-Specific Announce Policy Parameters B-41
BGP-3-Specific Announce Policy Parameters B-43
BGP-4-Specific Announce Policy Parameters B-47

Appendix C
Import and Export Route Filters

RIP C-1
RIP Import Filters C-1
RIP Export Filters C-5
OSPF Route Filters C-8
OSPF Import Filters C-8
OSPF Export Filters C-9
BGP-3 ..... C-13
BGP-3 Import Filters C-13
BGP-3 Export Filters C-17
EGP C-21
EGP Import Filters C-21
EGP Export Filters C-24

Appendix D
Route Weight Worksheet

Appendix E
IP/OSPF Configuration

Index

Figures

Figure 1-1. Network and Host Portions of IP Addresses 1-3
Figure 1-2. Internet Segmented into Three Autonomous Systems 1-9
Figure 1-3. IP Routing Table 1-14
Figure 1-4. Accept and Announce Policies 1-15
Figure 3-1. IP Interface 3-13
Figure 3-2. Multinet Configuration 3-15
Figure 3-3. IP Routers Source Routing across a Token Ring Network 3-23
Figure 3-4. ARP Example 3-27
Figure 3-5. Proxy ARP Example 3-31
Figure 3-6. RIPSO Security Label 3-39
Figure 3-7. RIPSO Network 3-44
Figure 3-8. Blacker Front End Network Configuration 3-48
Figure 5-1. OSPF Areas 5-5
Figure 5-2. Point-to-Multipoint Topology 5-16
Figure 5-3. Configurable Cost Metrics Usage Example 5-23
Figure 5-4. OSPF ASE Routes 5-27
Figure 5-5. AS External Route Tag 5-29
Figure 5-6. Area Border Router 5-31
Figure 5-7. Virtual Link and Transit Area 5-33
Figure 6-1. BGP Connecting Autonomous Systems Running OSPF 6-2
Figure 6-2. Transit Autonomous System 6-4
Figure 6-3. BGP over an Unnumbered Point-to-point Link 6-14
Figure 6-4. Establishing and Confirming a Connection between BGP Peers 6-16
Figure 7-1. EGP Connection between Two Autonomous Systems Running RIP 7-2
Figure 8-1. NetBIOS over IP 8-2
Figure 8-2. Broadcasting a Name Query Request 8-4
Figure 8-3. Returning a Unicast Name Query Response 8-5
Figure E-1. IP/OSPF Configuration E-2

Tables

Table 1-1. Subnet Masks for Class B and Class C Addresses 1-6
Table 1-2. IP Router RFC Support 1-17
Table 3-1. Source Route Bridge Support for Host-Only Mode 3-5
Table 3-2. Learning Bridge Support for Host-Only Mode 3-6
Table 3-3. BFE X.25 Packet-Level Parameter Settings 3-51
Table 3-4. BFE X.25 Network Service Record Parameter Settings 3-53
Table 6-1. BGP-3 Path Attributes 6-5
Table 6-2. BGP-4 Optional Path Attributes 6-6
Table 6-3. Characters in AS Path Pattern Matching 6-23
Table 7-1. Router Mode Determinator 7-3
Table E-1. Internal Backbone Router 1 ..... E-3
Table E-2. Area Border Router 2 ..... E-4
Table E-3. Area Border Router 3 ..... E-5
Table E-4. Area Border Router 4 ..... E-6
Table E-5. Internal Backbone Router 5 ..... E-7
Table E-6. AS Boundary Router 6 ..... E-7

About This Guide

If you are responsible for configuring and managing Bay Networks™ routers, you need to read this guide. This guide provides an overview of the Configuration Manager, and instructions about

- Starting up a router configuration file
- Configuring the router's connection to the Technician Interface console
- Specifying administrative information
- Implementing configurations created in local or remote mode, and saving dynamically made changes to a configuration file
- Using the Configuration Change utility
- Configuring a router with a new link module or net module

Before You Begin

Before using this guide, you must complete the following procedures. For a new router:

- Install the router (refer to the installation manual that came with your router).
- Connect the router to the network and create a pilot configuration file (refer to *Quick-Starting Routers and BNX Platforms*, *Connecting AN and ANH Systems to a Network*, or *Connecting ASN Routers to a Network)*.

Make sure that you are running the latest version of Bay Networks Site Manager and router software. For instructions, refer to *Upgrading Routers from Version 7–10*.xx *to Version 11.0*.

Conventions

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: if command syntax is **ping** *<ip_address>*, you enter **ping 192.32.10.12**
bold text	Indicates text that you need to enter, command names, and buttons in menu paths. Example: Enter **wfsm &**
	Example: Use the **dinfo** command.
	Example: ATM DXI > Interfaces > **PVCs** identifies the PVCs button in the window that appears when you select the Interfaces option from the ATM DXI menu.
brackets ([ ])	Indicate optional elements. You can choose none, one, or all of the options.
ellipsis points	Horizontal (. . .) and vertical ellipsis points indicate omitted information.
italic text	Indicates variable values in command syntax descriptions, new terms, file and directory names, and book titles.
quotation marks (“ ”)	Indicate the title of a chapter or section within a book.
`screen text`	Indicates data that appears on the screen. Example: `Set Bay Networks Trap Monitor Filters`
separator (>)	Separates menu and option names in instructions and internal pin-to-pin wire connections. Example: Protocols > AppleTalk identifies the AppleTalk option in the Protocols menu.
	Example: Pin 7 > 19 > 20
vertical line (|)	Indicates that you enter only one of the parts of the command. The vertical line separates choices. Do not type the vertical line when entering the command. Example: If the command syntax is
	show at routes | **nets**, you enter either **show at routes** or **show at nets**, but not both.

Acronyms

AUI	Attachment Unit Interface
BootP	Bootstrap Protocol
BRI	Basic Rate Interface
CCITT	International Telegraph and Telephone Consultative Committee (now ITU-T)
CSMA/CD	carrier sense multiple access with collision detection
DLCMI	Data Link Control Management Interface
GUI	graphical user interface
HDLC	high-level data link control
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
ITU-T	International Telecommunications Union–Telecommunications (formerly CCITT)
LAN	local area network
MAC	media access control
MAU	media access unit
MDI-X	media-dependent interface with crossover
NBMA	nonbroadcast multi-access
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First (Protocol)
PPP	Point-to-Point Protocol
RIP	Routing Information Protocol
SMDS	switched multimegabit data service
SNMP	Simple Network Management Protocol
STP	shielded twisted-pair
TCP/IP	Transmission Control Protocol/Internet Protocol
TELNET	Telecommunication Network
TFTP	Trivial File Transfer Protocol
TPE	twisted-pair Ethernet
UTP	unshielded twisted-pair
WAN	wide area network

Ordering Bay Networks Publications

To purchase additional copies of this document or other Bay Networks publications, order by part number from the Bay Networks Press™ at the following telephone or fax numbers:

- Telephone 1-800-845-9523
- Fax (United States/Canada) 1-800-582-8000
- Fax (International) 1-916-939-1010

You can also use these numbers to request a free catalog of Bay Networks Press product publications.

Technical Support and Online Services

To ensure comprehensive network support to our customers and partners worldwide, Bay Networks Customer Service has Technical Response Centers in key locations around the globe:

- Billerica, Massachusetts
- Santa Clara, California
- Sydney, Australia
- Tokyo, Japan
- Valbonne, France

The Technical Response Centers are connected via a redundant Frame Relay Network to a Common Problem Resolution system, enabling them to transmit and share information, and to provide live, around-the-clock support 365 days a year.

Bay Networks Information Services complement the Bay Networks Service program portfolio by giving customers and partners access to the most current technical and support information through a choice of access/retrieval means. These include the World Wide Web, CompuServe, Support Source CD, Customer Support FTP, and InfoFACTS document fax service.

Bay Networks Customer Service

If you purchased your Bay Networks product from a distributor or authorized reseller, contact that distributor's or reseller's technical support staff for assistance with installation, configuration, troubleshooting, or integration issues.

Customers can also purchase direct support from Bay Networks through a variety of service programs. As part of our PhonePlus™ program, Bay Networks Service sets the industry standard, with 24-hour, 7-days-a-week telephone support available worldwide at no extra cost. Our complete range of contract and noncontract services also includes equipment staging and integration, installation support, on-site services, and replacement parts delivery -- within approximately 4 hours.

To purchase any of the Bay Networks support programs, or if you have questions on program features, use the following numbers:

Region	Telephone Number	Fax Number
United States and Canada	1-800-2LANWAN; enter Express Routing Code (ERC) 290 when prompted (508) 436-8880 (direct)	(508) 670-8766
Europe	(33) 92-968-300	(33) 92-968-301
Asia/Pacific Region	(612) 9927-8800	(612) 9927-8811
Latin America	(407) 997-1713	(407) 997-1714

In addition, you can receive information on support programs from your local Bay Networks field sales office, or purchase Bay Networks support directly from your authorized partner.

Bay Networks Information Services

Bay Networks Information Services provide up-to-date support information as a first-line resource for network administration, expansion, and maintenance. This information is available from a variety of sources.

World Wide Web

The Bay Networks Customer Support Web Server offers a diverse library of technical documents, software agents, and other important technical information to Bay Networks customers and partners.

A special benefit for contracted customers and resellers is the ability to access the Web Server to perform Case Management. This feature enables your support staff to interact directly with the network experts in our worldwide Technical Response Centers. A registered contact with a valid Site ID can

- View a listing of support cases and determine the current status of any open case. Case history data includes severity designation, and telephone, e-mail, or other logs associated with the case.
- Customize the listing of cases according to a variety of criteria, including date, severity, status, and case ID.
- Log notes to existing open cases.
- Create new cases for rapid, efficient handling of noncritical network situations.
- Communicate directly via e-mail with the specific technical resources assigned to your case.

The Bay Networks URL is *http://www.baynetworks.com*. Customer Service is a menu item on that home page.

Customer Service FTP

Accessible via URL *ftp://support.baynetworks.com* (134.177.3.26), this site combines and organizes support files and documentation from across the Bay Networks product suite, including switching products from our Centillion™ and Xylogics® business units. Central management and sponsorship of this FTP site lets you quickly locate information on any of your Bay Networks products.

Support Source CD

This CD-ROM -- sent quarterly to all contracted customers -- is a complete Bay Networks Service troubleshooting knowledge database with an intelligent text search engine.

The Support Source CD contains extracts from our problem-tracking database; information from the Bay Networks Forum on CompuServe; comprehensive technical documentation, such as Customer Support Bulletins, Release Notes, software patches and fixes; and complete information on all Bay Networks Service programs.

You can run a single version on Macintosh Windows 3.1, Windows 95, Windows NT, DOS, or UNIX computing platforms. A Web links feature enables you to go directly from the CD to various Bay Networks Web pages.

CompuServe

For assistance with noncritical network support issues, Bay Networks Information Services maintain an active forum on CompuServe, a global bulletin-board system. This forum provides file services, technology conferences, and a message section to get assistance from other users.

The message section is monitored by Bay Networks engineers, who provide assistance wherever possible. Customers and resellers holding Bay Networks service contracts also have access to special libraries for advanced levels of support documentation and software. To take advantage of CompuServe's recently enhanced menu options, the Bay Networks Forum has been re-engineered to allow links to our Web sites and FTP sites.

We recommend the use of CompuServe Information Manager software to access these Bay Networks Information Services resources. To open an account and receive a local dial-up number in the United States, call CompuServe at 1-800-524-3388. Outside the United States, call 1-614-529-1349, or your nearest CompuServe office. Ask for Representative No. 591. When you are on line with your CompuServe account, you can reach us with the command **GO BAYNET**.

InfoFACTS

InfoFACTS is the Bay Networks free 24-hour fax-on-demand service. This automated system has libraries of technical and product documents designed to help you manage and troubleshoot your Bay Networks products. The system responds to a fax from the caller or to a third party within minutes of being accessed.

To use InfoFACTS in the United States or Canada, call toll-free 1-800-786-3228. Outside North America, toll calls can be made to 1-408-764-1002. In Europe, toll-free numbers are also available for contacting both InfoFACTS and CompuServe. Please check our Web page for the listing in your country.

How to Get Help

Use the following numbers to reach your Bay Networks Technical Response Center:

Technical Response Center	Telephone Number	Fax Number
Billerica, MA	1-800-2LANWAN	(508) 670-8765
Santa Clara, CA	1-800-2LANWAN	(408) 764-1188
Valbonne, France	(33) 92-968-968	(33) 92-966-998
Sydney, Australia	(612) 9927-8800	(612) 9927-8811
Tokyo, Japan	(81) 3-5402-0180	(81) 3-5402-0173

Chapter 1
IP Concepts, Terminology, and Features

The following sections introduce concepts and terminology used in this manual:

- IP Addresses -- page 1-2
  - Subnet Addressing -- page 1-4
  - Supernet Addressing -- page 1-7
  - Classless Interdomain Routing -- page 1-8
- Autonomous Systems -- page 1-8
- Routing Protocols -- page 1-8
  - Routing Information Protocol (RIP) -- page 1-9
  - Open Shortest Path First (OSPF) Protocol -- page 1-10
  - Border Gateway Protocol (BGP) -- page 1-10
  - Exterior Gateway Protocol (EGP) -- page 1-10
- Route Preferences -- page 1-11
- Route Weights -- page 1-11
- IP Routing Policies and Filters -- page 1-13
- IP Traffic Filters -- page 1-17
- RFC Compliance -- page 1-17

IP Addresses

An IP address consists of 32 bits that have the form *network.host*. The *network* portion is a network number ranging from 8 to 24 bits. The *host* portion is the remaining 8 to 24 bits identifying a specific host on the network. The Internet Network Information Center (NIC) assigns the network portion of the IP address. Your network administrator assigns the host portion.

NIC recognizes three primary classes of networks: A, B, and C. In addition, NIC has recently identified two other classes: Class D for networks that support multicasting, which allows an IP datagram to be transmitted to a single multicast group consisting of hosts spread across separate physical networks; and Class E for experimental networks. The IP router does not fully support Class D or Class E networks.

Based on the size of the network, the NIC classifies a network as Class A, B, or C (the most common). The network class determines the number of bits assigned to the network and host portions of the IP address, as follows:

Network Size	**Class**	**Network Portion**	**Host Portion**
More than 65,534 hosts	A	8 bits	24 bits
254 to 65,533 hosts	B	16 bits	16 bits
Fewer than 254 hosts	C	24 bits	8 bits

The position of the first bit set to 0 (whether it is the first, second, third, or fourth bit) in the first octet of an IP address indicates the network Class (A, B, C, or D). If no bit is set to 0, it is a Class E network. Figure 1-1 shows the placement of the first bit set to 0 for Class A, B, and C networks. The figure also shows how a network's class affects the network and host portions of the IP address.

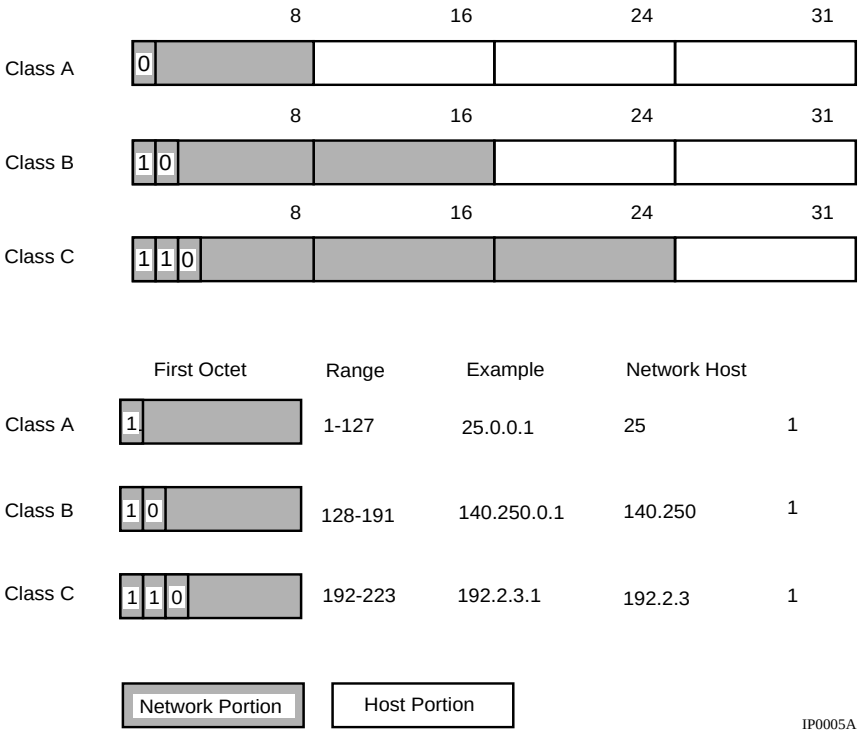

Figure 1-1. Network and Host Portions of IP Addresses

You specify IP addresses in dotted decimal notation. To express an IP address in dotted decimal notation, you convert each 8-bit octet of the IP address to a decimal number and separate the numbers by decimal points.

For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted decimal notation as 128.32.10.167. The most significant 2 bits (10) in the first octet indicate that the network is Class B; therefore, the first 16 bits compose the NIC-assigned network portion field. The third octet (00001010) and fourth octet (10100111) compose the host field.

Subnet Addressing

The concept of subnetworks (or subnets) extends the IP addressing scheme. Subnets are two or more physical networks that share a common network-identification field (the NIC-assigned network portion of the 32-bit IP address). Subnets allow an IP router to hide the complexity of multiple LANs from the rest of the internet.

With subnets, you partition the host portion of an IP address into a subnet number and a “real” host number on that subnet. The IP address is then defined by *network.subnet.host*. Routers outside the network do not interpret the subnet and host portions of the IP address separately.

Routers inside a network containing subnets use a 32-bit subnet mask that identifies the extension bits. In *network.subnet.host*, the *subnet.host* portion (or the local portion) contains an arbitrary number of bits. The network administrator allocates bits within the local portion to subnet and host, and then assigns values to subnet and host.

For example, the following is the IP address of a network that contains subnets: 10000000 00100000 00001010 10100111. You specify this address in dotted decimal notation as 128.32.10.167.

The second bit of the first octet is set to 0, indicating that the network is a Class B network. Therefore, the NIC-assigned network portion contains 16 bits, and the locally assigned local portion contains 16 bits.

The network administrator allocates the 16 bits in the local portion field as follows:

- Allocates the upper 8 bits (00001010) with a value of 10 to the subnet portion
- Allocates the lower 8 bits (10100111) with a value of 167 to the host portion

In other words, the 16-bit local portion field, together with the 16-bit network field, specify host 167 on Subnet 10 of network 128.32.

You now need a subnet mask to identify those bits in the 32-bit IP address that specify the network field and those bits that specify the subnet field. Like the IP address, you specify the subnet mask in dotted decimal notation.

You construct a subnet mask as follows:

- Assign a value of 1 to each of the 8, 16, or 24 bits in the network field.

- Assign a value of 1 to each bit in the subnet field.
- Assign a value of 0 to each bit in the host field.
- Convert the resulting 32-bit string to dotted decimal notation.

For example, to construct a subnet mask for the IP address described earlier (10000000 00100000 00001010 10100111), do the following:

1. Assign a value of 1 to each bit in the network field.

   The position of the first bit set to 0 in the first octet of the IP address indicates that the network is Class B; therefore, the network field contains 16 bits: 11111111 11111111.

2. Assign a value of 1 to each bit in the subnet field.

   The network administrator allocated the upper 8 bits of the local portion to the subnet portion, as follows: 11111111.

3. Assign a value of 0 to each bit in the host field.

   The network administrator allocated the lower 8 bits of the local portion field to the host identification, as follows: 00000000.

4. Convert the resulting 32-bit string (11111111 11111111 11111111 00000000) to dotted decimal notation, as follows: 255.255.255.000.

Table 1-1 shows the range of possible subnet masks for Class B and Class C addresses, along with the number of bits that the mask allocates for a subnet address, the number of recommended subnets associated with the mask, and the number of hosts per subnet.

Table 1-1. Subnet Masks for Class B and Class C Addresses

Number of Bits	Subnet Mask	Number of Subnets (Recommended)	Number of Hosts per Subnet
Class B			
2	255.255.192.0	2	16,382
3	255.255.224.0	6	8,190
4	255.255.240.0	14	4,094
5	255.255.248.0	30	2,046
6	255.255.252.0	62	1,022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1,022	62
11	255.255.255.224	2,046	30
12	255.255.255.240	4,094	14
13	255.255.255.248	8,190	6
14	255.255.255.252	16,382	2
Class C			
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

Supernet Addressing

A supernet is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the following block of contiguous 32-bit addresses (192.32.0.0 through 192.32.7.0 in decimal notation).

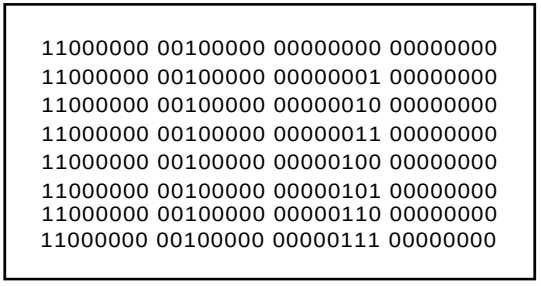

IP0007A

The supernet address for this block is 11000000 00100000 00000, the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an *address*/*mask* pair:

- *address* is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in decimal notation).
- *mask* is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted decimal notation).

The complete supernet address in this example is 192.32.0.0/255.255.248.0.

Classless Interdomain Routing

Classless interdomain routing (CIDR) is an addressing scheme that employs supernet addresses to represent multiple IP destinations. Rather than advertise a separate route for each destination in a supernet, a router can use a supernet address to advertise a single route -- called an *aggregate* route -- that represents all of the destinations. This reduces the size of the routing tables used to store advertised IP routes.

BGP-4 supports classless interdomain routing. OSPF supports classless routing within a domain.

Autonomous Systems

LANs and WANs interconnected by IP routers form a group of networks called an *internet*. For administrative purposes, an internet is divided into autonomous systems. An *autonomous system* (AS) is simply a collection of routers (called gateways in IP terminology) and hosts. Figure 1-2 depicts a sample internet segmented into three autonomous systems.

Routing Protocols

Routers inside an autonomous system use an *interior gateway protocol* to communicate network topology changes to each other. Routers in separate autonomous systems use an *exterior gateway protocol* to communicate. The IP router implements two dynamic interior protocols: RIP and OSPF. The IP router implements two exterior protocols: BGP and EGP.

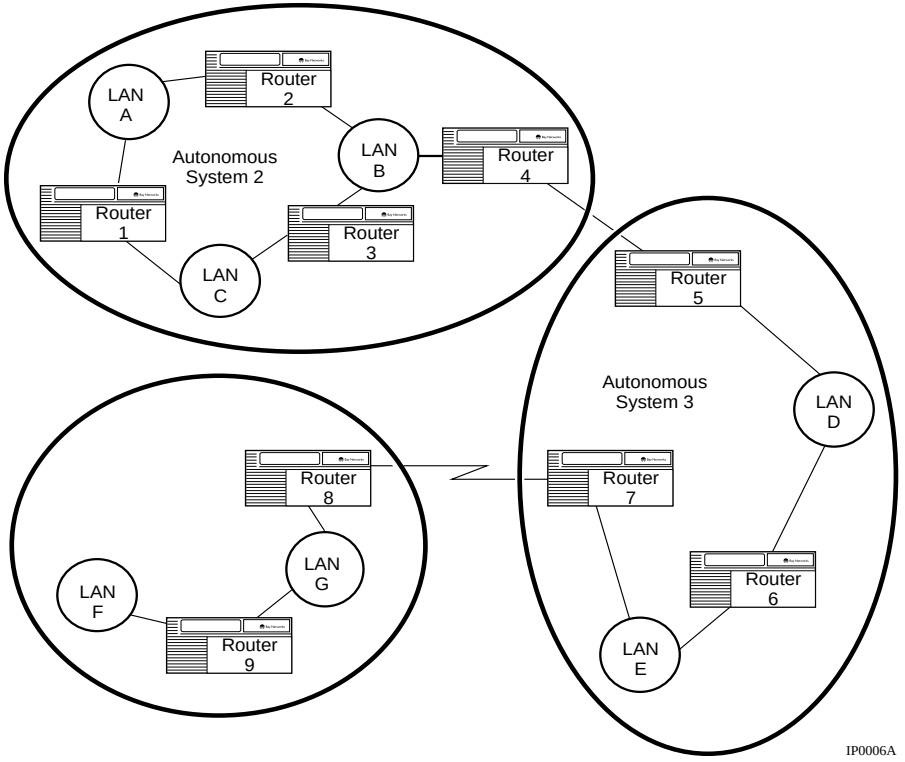

Figure 1-2. Internet Segmented into Three Autonomous Systems

Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is a distance-vector protocol that enables routers in the same autonomous system to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring networks and listen for RIP updates from the routers on those neighboring networks. Routers use the information in the RIP updates to keep their internal routing tables current. For RIP, the “best” path to a destination is the shortest path (the path with the fewest hops). RIP computes distance as a metric, usually the number of hops (or routers) from the origin network to the target network.

Open Shortest Path First (OSPF) Protocol

The Open Shortest Path First (OSPF) protocol is an interior gateway protocol (IGP) intended for use in large networks. Using a link state algorithm, OSPF exchanges routing information between routers in an autonomous system. Routers synchronize their topological databases. Once the routers are synchronized and the routing tables are built, the routers will flood topology information only in response to some topological change. For OSPF, the “best” path to a destination is the path that offers the least cost metric delay. In OSPF, cost metrics are configurable, allowing you to specify preferred paths.

OSPF supports CIDR and can carry supernet advertisements within a routing domain.

Border Gateway Protocol (BGP)

The Border Gateway Protocol (BGP) is an exterior gateway protocol used to exchange network reachability information with other BGP systems. BGP routers form relationships with other BGP routers. Using an entity called a BGP speaker, BGP routers transmit and receive current routing information over a reliable transport layer connection. Because a reliable transport mechanism is used, periodic updates are not necessary.

BGP updates contain “path attributes” that describe the route to a set of destination networks. When multiple paths are available, BGP compares these path attributes to choose the preferred path.

BGP-3 and BGP-4 are supported. BGP-4 is the border gateway protocol that supports CIDR.

Exterior Gateway Protocol (EGP)

The Exterior Gateway Protocol (EGP-2) is an exterior gateway protocol used to exchange network reachability information between routers in different autonomous systems. An IGP, such as RIP or OSPF, is used within an AS to facilitate the communication of routing information within the autonomous system. The routers that serve as the end points of a connection between two autonomous systems run an exterior gateway protocol, such as EGP-2.

Routers establish EGP neighbor relationships in order to periodically exchange reliable network reachability information. The router uses this information to maintain a list of gateways, the networks the gateways can reach, and the corresponding distances.

Route Preferences

The IP router maintains an internal routing table. When determining how to forward a datagram, the IP router consults the table to determine the specific route a datagram should take. A routing table can contain direct routes for the IP router's network interfaces, static routes, and the routes learned from RIP, OSPF, BGP, and/or EGP, if enabled (information about adjacent hosts is maintained in a separate table).

It is possible for a routing table to contain multiple routes to the same destination. In such a situation, IP uses (among other information) a preference value to determine which route to select. Preference values range from 1 to 16 (the higher the number, the greater the preference).

By default, RIP, BGP, EGP, and OSPF external routes have a preference value of 1. Static routes, direct routes, and OSPF intra-area and interarea routes have a default preference of 16.

You can configure a preference value in the range of 1 to 16 for RIP, BGP, EGP, OSFP external, and static routes. The preference of direct routes and OSPF intra-area and interarea routes cannot be user-configured.

To assign a preference to a route learned by RIP, OSPF, BGP, and EGP, you configure an accept policy for the route. If an incoming route matches the policy, IP assigns the preference value you specify to the route and considers the route for possible inclusion in the routing table.

Route Weights

Route-weight calculation is an internal tool that IP uses to facilitate selection of the best route among alternative routes to the same destination. Route-selection criteria are encoded into the route weight in a way that allows IP to compare routes simply by comparing their weight values, regardless of route sources.

Appendix D contains a worksheet that you can you use to calculate route weights in your configuration.

Route-weight calculation increases the efficiency of the route-selection process and at the same time reduces the size of the routing database, since all route selection parameters for each route are encoded in a single integer -- the weight value -- rather than stored in separate variables.

Using selection criteria encoded in the route weight, IP chooses routes in the following order:

- The route with the highest preference value (see “Route Preferences” on page 1-11)
- A direct or OSPF intra-area route with the lowest metric
- A direct route with the lowest metric
- An OSPF intra-area route with the lowest metric
- An OSPF interarea route with the lowest metric
- An OSPF Type 1 external route with the lowest metric
- A BGP route with the highest LOCAL_PREF value
- A RIP route with the lowest metric
- An EGP route with the lowest metric
- A static route with the lowest metric
- An OSPF Type 2 external route with a pre-Version-8.00-style metric

Note: If OSPF is configured to propagate external routes using the route weight as the Type 2 metric, routes that are received as OSPF ASE Type 2 routes are evaluated according to their respective origins (for example, RIP or BGP).

IP Routing Policies and Filters

The IP router allows you to control the flow of routing data to and from the routing tables. This control is provided by two mechanisms:

- IP accept and announce policies
- IP import and export filters

Note: Accept and announce policies provide a superset of the parameters provided by import and export filters. We currently support both IP policies and IP route filters. However, network administrators using import and export filters for routing table management should migrate as quickly as possible to IP policies. In a future release, support for import and export filters will be discontinued.

IP *accept policies* (and the subset of parameters provided by import filters) govern the addition of new RIP-, OSPF-, BGP-, or EGP-derived routes to the routing tables. When RIP, OSPF, BGP, or EGP receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

IP *announce policies* (and the subset of parameters provided by export filters) govern the propagation of RIP, OSPF, BGP, or EGP routing information. When preparing a routing advertisement, RIP, OSPF, BGP, or EGP consults its announce policies to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure).

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the Internet protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols. Figure 1-3 shows a router configured with all of the Internet protocols supported by Bay Networks: OSPF, RIP, BGP-3, BGP-4, and EGP. The arrows indicate the direction of flow of routing information between the network and the protocols running on the router, between the protocols and the routing table manager, and between the routing table manager and the routing table.

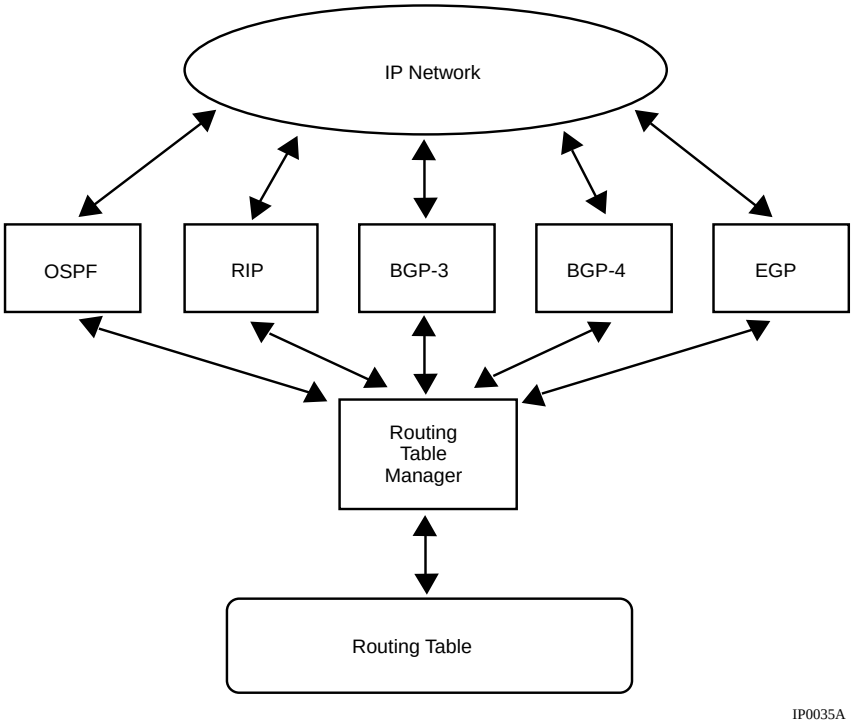

Figure 1-3. IP Routing Table

The flow of routing information between the network, the protocols, and the routing table manager is controlled by *routing information policies.*

Each time a routing update arrives from a remote router, the following steps occur (see Figure 1-4):

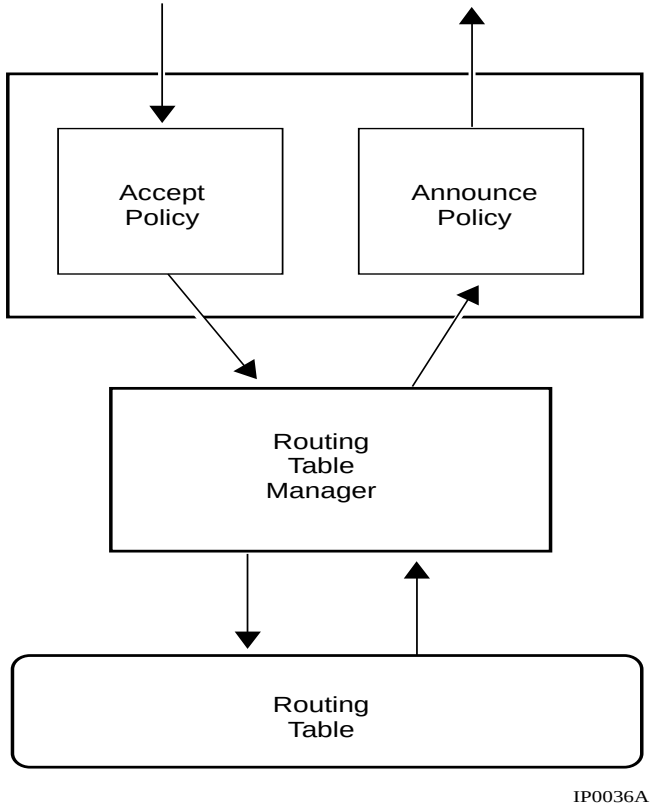

Figure 1-4. Accept and Announce Policies

1. The protocol receiving the route consults an *accept policy* to determine whether to forward the route to the IP routing table manager or drop the route.
2. If the protocol forwards the route, the routing table manager determines whether to inject the route into the routing table.

Periodically, the routing table manager announces routes to other routers in the network:

1. The routing table manager forwards a route for advertisement to the protocol.

2. The protocol consults an *announce policy* to determine whether or not to advertise the route to the network.

Note: The way OSPF applies accept and announce policies to routing information differs in several ways from the procedure shown in Figure 1-4. OSPF link-state advertisements (LSAs) are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them on to other routers. To accomplish this, OSPF accept and announce policies act in the following manner:

OSPF accept policies control which OSPF non-self-originated external routing information is passed to the routing table manager. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.

OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router's self-originated information.

IP accept and announce policies and policy parameters are described in Appendix B.

IP import and export filters and filter parameters are described in Appendix C.

IP Traffic Filters

A traffic filter enables the router to selectively relay or drop an inbound packet, frame, or datagram based on standard protocol fields or user-defined fields. Traffic filters apply to incoming traffic only.

For information on IP traffic filters, see *Configuring Traffic Filters and Protocol Prioritization.*

RFC Compliance

Table 1-2 lists the Internet Requests for Comments (RFCs) with which the IP router complies. This manual assumes you are familiar with these RFCs.

Table 1-2. IP Router RFC Support

RFC	Specifies
768	User Datagram Protocol (UDP)
783	Trivial File Transfer Protocol (TFTP)
791	Internet Protocol (IP)
792	Internet Control Message Protocol (ICMP)
793 and 1323	Transmission Control Protocol
826	Address Resolution Protocol (ARP)
903	RARP Server
904	EGP-2
950	Internet subnetting procedures
951	BOOTP
1009	Internet Gateways
1027	Proxy ARP
1042	IP over IEEE 802.*x* networks
1058 and 1388	Routing Information Protocol (RIP)
1063	Maximum Transmission Unit (MTU) discovery option
1108	RIPSO
1112	Host Extensions for IP Multicasting
1157	Simple Network Management Protocol (SNMP)

(continued)

Table 1-2. IP Router RFC Support *(continued)*

RFC	Specifies
1188	IP over FDDI networks
1209	IP over SDMS
1256	ICMP Router Discovery Messages
1267	BGP-3
1293	Inverse ARP over Frame Relay
1332	IP over PPP
1356	IP over X.25
1403	BGP OSPF Interaction
1483	IP over ATM DXI, IP over PVC, IP Multicast over PVC
1490	IP over Frame Relay
1577	IP over SVC
1583	Open Shortest Path First (OSPF) Protocol Version 2
1771	BGP-4

Chapter 2
Starting IP Services

The following sections show you how to perform a basic configuration -- that is, a configuration using all available defaults -- for the IP services described in this manual:

- Starting IP -- page 2-3
  - Deleting IP from an Interface -- page 2-4
  - Customizing IP -- page 2-4
- Starting RIP -- page 2-5
  - Adding RIP to an IP Interface -- page 2-5
  - Deleting RIP from An IP Interface -- page 2-6
  - Customizing RIP -- page 2-6
- Starting OSPF -- page 2-7
  - Adding OSPF to an IP Interface -- page 2-8
  - Deleting OSPF from an IP Interface -- page 2-8
  - Customizing OSPF -- page 2-8
- Starting BGP -- page 2-9
  - Deleting BGP from the Router -- page 2-10
  - Deleting BGP-3 and BGP-4 from the Router -- page 2-11
  - Customizing BGP -- page 2-11
- Starting EGP -- page 2-12
  - Deleting EGP from the Router -- page 2-12
  - Customizing EGP -- page 2-13

- Starting NetBIOS over IP -- page 2-13

  Adding NetBIOS to an IP Interface -- page 2-13
  Deleting NetBIOS from an IP Interface -- page 2-14
  Customizing NetBIOS over IP -- page 2-14

- Using the Circuitless IP Interface -- page 2-15

  Starting IP on the Circuitless Interface -- page 2-15
  Adding OSPF and BGP to the Circuitless Interface -- page 2-16
  Selecting Slots to Support the Circuitless Interface -- page 2-16

- Configuring an Unnumbered IP Interface -- page 2-16

  Using the Alternate Associated Address Option -- page 2-18

Starting IP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as an interface to an attached network. For information and instructions, see *Configuring Line Services*.

When you have successfully configured the circuit, the Select Protocols window opens.

1. **Select IP from the Select Protocols window and click on OK.**

   The IP Configuration window opens.

2. **Supply the following information:**

   - *An IP address for this interface.* There is no default for the IP address. You must supply an address or enter 0.0.0.0 to indicate that this is an unnumbered interface. (For information about unnumbered interfaces, see "Configuring an Unnumbered IP Interface" on page 2-16.)

     Site Manager: IP Address parameter: page A-22

   - *A subnet mask.* When you specify an address, IP automatically calculates an appropriate subnet mask, depending on the class of network. You can use Site Manager to override the default mask.

     Site Manager: Subnet Mask parameter: page A-22

   - *A transmit broadcast address*. By default, IP uses a broadcast address that consists of all 1s. You can use Site Manager to override the default.

     Site Manager: Transmit Bcast Addr parameter: page A-23

   - *An associated IP address.* If you are configuring an unnumbered IP interface, use Site Manager to supply an associated IP address.(For information about unnumbered interfaces, see "Configuring an Unnumbered IP Interface" on page 2-16.)

     Site Manager: UnNumbered Assoc Address parameter: page A-23

3. **Click on OK.**

Deleting IP from an Interface

To delete IP from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector from which you want to delete IP services.**
2. **Click on Edit Circuit.**
3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens. The IP button is highlighted to show that IP is enabled on the circuit.
4. **Click on IP to delete it.**
5. **Click on OK to exit the window.**
6. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Customizing IP

The instructions in this chapter show you how to start IP using all default values and settings.

For information on modifying IP settings, see Chapter 3.

Starting RIP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring Line Services*.

The Select Protocols window opens.

1. **Select IP and RIP from the Select Protocols window and click on OK.**

   Site Manager displays the IP Configuration window.

2. **Edit the IP Configuration window as described in "Starting IP" on page 2-3.**

3. **Click on OK.**

Adding RIP to an IP Interface

To add RIP to an IP interface, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector to which you want to add RIP services.**

2. **Click on Edit Circuit.**

3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens.

4. **Click on RIP to select it.**

5. **Click on OK to exit the window.**

6. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Deleting RIP from An IP Interface

To delete RIP from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector from which you want to delete RIP services.**
2. **Click on Edit Circuit.**
3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens. The RIP button is highlighted to show that RIP is enabled on the IP interface.
4. **Click on RIP to delete it.**
5. **Click on OK to exit the window.**
6. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Customizing RIP

The instructions in this chapter show you how to start RIP using all default values and settings.

For information on modifying RIP defaults, see Chapter 4.

Starting OSPF

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring Line Services*.

The Select Protocols window opens.

1. **Select IP and OSPF from the Select Protocols window and click on OK.**

   Site Manager displays the IP Configuration window.

2. **Edit the parameters on the IP Configuration window as described in "Starting IP" on page 2-3.**

   Site Manager displays the Initial OSPF Global Configuration window.

   This window contains a subset of the parameters that you can set to configure the OSPF soloist.

3. **When you are finished with the Initial Global Configuration window, click on OK.**

   An interface-specific OSPF Area Configuration window opens.

4. **Enter the ID of the area to which this interface belongs.**

   By default, OSPF assumes that the interface is connected to a network in the backbone (area 0.0.0.0). Use Site Manager to specify an area ID.

Site Manager: Area ID parameter: page A-78

5. **Click on OK.**

   The Network Type window opens.

6. **Specify the type of network to which this interface is connected.**

   By default, OSPF assumes that the interface is connected to a broadcast network. Use Site Manager to specify the network type.

Site Manager: Type parameter: page A-79

7. **Click on OK.**

Adding OSPF to an IP Interface

To add OSPF to an IP interface, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector to which you want to add OSPF services.**
2. **Click on Edit Circuit.**
3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens.
4. **Click on OSPF to select it.**
5. **Click on OK to exit the window.**
6. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Deleting OSPF from an IP Interface

To delete OSPF from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector from which you want to delete OSPF services.**
2. **Click on Edit Circuit.**
3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens. The OSPF button is highlighted to show that OSPF is enabled on the circuit.
4. **Click on OSPF to delete it.**
5. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Customizing OSPF

The instructions in this chapter show you how to start OSPF using all default values and settings.

For information on modifying OSPF defaults, see Chapter 5.

Starting BGP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring Line Services*.

The Select Protocols window opens.

1. **Select IP and BGP from the Select Protocols window and click on OK.**

   Site Manager displays the IP configuration window.

2. **Edit the IP Configuration Window as described in "Starting IP" on page 2-3.**

3. **Click on OK.**

   The BGP Configuration window opens.

4. **Supply the following information:**

   - *BGP identifier.* There is no default for this ID. Use Site Manager to supply the IP address of one of the router's interfaces.

     Site Manager: Identifier parameter: page A-1

   - *Local autonomous system number.* There is no default for this parameter. Use Site Manager to supply the appropriate AS number.

     Site Manager: Local AS parameter: page A-1

5. **Click on OK.**

   The BGP Peer window opens.

6. **Supply the following information:**

   - *IP address of the remote peer.* There is no default for this parameter. Use Site Manager to supply the IP address on the remote side of this peer-to-peer connection.

     Site Manager: Peer Address parameter: page A-6

   - *The remote peer AS number.* There is no default for this parameter. Use Site Manager to supply the number of the AS in which the remote peer is located.

     Site Manager: Peer AS parameter: page A-6

   - *IP address of the local peer.* There is no default for this parameter. Use Site Manager to supply the IP address of the interface on the local side of this peer-to-peer connection.

     Site Manager: Local Address parameter: page A-7

7. **Click on OK to enable default BGP service.**

Deleting BGP from the Router

You can delete BGP from all router circuits on which it is currently enabled.

To delete BGP, begin at the Configuration Manager window and complete the following steps:

1. **Select Protocols > IP > BGP > Delete BGP.**

   A pop-up window opens, prompting, “Do you really want to delete BGP?”

2. **Click on OK.**

   You are returned to the Configuration Manager window. BGP is removed from all circuits on the router.

Deleting BGP-3 and BGP-4 from the Router

You can delete BGP-3 from all router circuits on which it is currently enabled. To delete BGP-3, begin at the Configuration Manager window and complete the following steps:

1. **Select Protocols > IP > BGP > Delete BGP-3.**

   A popup window opens, prompting, “Do you really want to delete BGP-3?”

2. **Click on OK.**

   You are returned to the Configuration Manager window. BGP-3 is removed from all circuits on the router.

You can delete BGP-4 from all router circuits on which it is currently enabled. To delete BGP-4, begin at the Configuration Manager window and complete the following steps:

1. **Select Protocols > IP > BGP > Delete BGP-4.**

   A popup window opens, prompting, “Do you really want to delete BGP-4?”

2. **Click on OK.**

   You are returned to the Configuration Manager window. BGP-4 is removed from all circuits on the router.

Customizing BGP

The instructions in this chapter show you how to start BGP using all default values and settings.

For information on modifying BGP defaults, see Chapter 6.

Starting EGP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring Line Services*.

The Select Protocols window opens.

1. **Select IP and EGP from the Select Protocols window and click on OK.**

   Site Manager displays the IP configuration window.

2. **Edit the IP Configuration window as described in "Starting IP" on page 2-3.**

3. **Click on OK.**

   The EGP Neighbors Parameters window opens.

4. **Supply the following information:**

   - *IP address of the remote neighbor.* There is no default for this parameter. Use Site Manager to supply the address.

     Site Manager: Remote Autonomous System IP Address parameter: page A-19

   - *Gateway mode of the remote neighbor.* There is no default for this parameter. Use Site Manager to specify the mode: core or noncore.

     Site Manager: Gateway Mode parameter: page A-19

5. **Click on OK.**

Deleting EGP from the Router

You can delete EGP from all router circuits on which it is currently enabled. To delete EGP, begin at the Configuration Manager window and complete the following steps.

1. **Select Protocols > IP > BGP > Delete EGP.**

   A window opens, prompting, "Do you really want to delete EGP?"

2. **Click on OK.**

Customizing EGP

The instructions in this chapter show you how to start EGP using all default values and settings.

For information on modifying EGP defaults, see Chapter 7.

Starting NetBIOS over IP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring Line Services*.

The Select Protocols window opens.

1. **Select IP and NetBIOS from the Select Protocols window and click on OK.**

   Site Manager displays the IP configuration window.

2. **Edit the IP Configuration Window as described in "Starting IP" on page 2-3.**

3. **Click on OK.**

Adding NetBIOS to an IP Interface

To add NetBIOS to an IP interface, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector to which you want to add NetBIOS services.**

2. **Click on Edit Circuit.**

3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens.

4. **Click on NetBIOS to select it.**

5. **Click on OK to exit the window.**

6. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Deleting NetBIOS from an IP Interface

To delete NetBIOS from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. **Click on the connector from which you want to delete OSPF services.**
2. **Click on Edit Circuit.**
3. **Select Protocols > Add or Delete.**

   The Select Protocols window opens. The NetBIOS button is highlighted to show that NetBIOS is enabled on the circuit.
4. **Click on NetBIOS to delete it.**
5. **Select File > Exit to exit the Circuit Definition window and return to the Configuration Manager window.**

Customizing NetBIOS over IP

The instructions in this chapter show you how to start NetBIOS using all default values and settings.

For information on modifying NetBIOS defaults, see Chapter 8.

Using the Circuitless IP Interface

A *circuitless* IP interface has an IP address that is not mapped to a specific circuit. If one or more of the router's IP interfaces become disabled, this circuitless feature ensures that the router is always reachable using the circuitless IP interface address, as long as a viable path to the router exists. The IP router can support one circuitless IP interface.

IP traffic is delivered to and transmitted from the circuitless interface in the same way as any other IP interface. In addition, the circuitless IP interface can receive packets from any application.

When you configure a circuitless IP interface, note the following:

- You can configure one circuitless IP interface per router. Additional circuitless IP interfaces will not initialize.
- You can add IP routing protocols to the circuitless interface. (The current version of Site Manager allows you to add BGP and OSPF to a circuitless interface.)
- You must assign a unique IP address and subnetwork number to the circuitless IP interface.
- You *cannot* configure a circuitless IP interface in nonforwarding mode.

Starting IP on the Circuitless Interface

To configure a circuitless IP interface, begin at the Configuration Manager window and proceed as follows:

1. **Select Protocols > IP > Circuitless IP > Create to display the IP Configuration window.**
2. **Edit the parameters on the IP Configuration window.**
3. **Click on OK to save the circuitless IP interface.**

   A special Select Protocols window opens, listing the protocols you can configure on a circuitless interface.
4. **Select a protocol and click on OK.**

Adding OSPF and BGP to the Circuitless Interface

Once you have created a circuitless IP interface on the router, you can add OSPF and BGP to the interface. Begin at the Configuration Manager window and proceed as follows:

1. **Select Protocols > IP > Circuitless IP > Change Protocols.**

   A Select Protocols window opens, listing the protocols you can add to the circuitless interface.

2. **Select a protocol and click on OK.**

   The router adds the protocol to the circuitless IP interface.

Selecting Slots to Support the Circuitless Interface

By default, all slots support the IP circuitless interface.

You can use Site Manager to specify the slots that can support the circuitless interface.

Site Manager: Slot Mask parameter: page A-34

Configuring an Unnumbered IP Interface

IP allows you to configure an interface on a point-to-point connection without using an IP address. Such an interface is called an *unnumbered interface*. Point-to-point connections using unnumbered interfaces can be configured to advertise RIP (see Note), OSPF, IBGP, DVMRP, and static routes.

The ability to establish a point-to-point link using an unnumbered IP interface helps alleviate two of the major problems caused by the continued rapid growth of the Internet: exhaustion of Class B network addresses and exhaustion of the 32-bit IP address space.

The network administrator associates each unnumbered interface with the IP address of any numbered interface on the router, including the circuitless interface. The router can support multiple unnumbered interfaces. Multiple unnumbered interfaces can be associated with the same IP address.

Since all traffic over an unnumbered interface uses broadcast addressing at the link layer, neither an adjacent host specification or address resolution is required.

Note: The associated address assigned to the unnumbered interface determines whether or not RIP configured to send updates in V1 mode will advertise a subnetwork over the unnumbered interface. The associated address also determines which mask is applied to RIP V1 updates received on that interface. For unnumbered links using RIP V1, the defined associated addresses at each end of the link must belong to the same network and have the same mask for routes to be exchanged correctly.

If a subnetwork on the router has the same mask as the associated address, RIP V1 will advertise that subnet over the unnumbered interface. If the mask on the subnetwork is different from the mask of the associated address, RIP V1 advertises only the natural network of the subnet.

We recommend that you select RIP2 mode for unnumbered interfaces. With RIP2, RIP updates contain both the route and mask information.

Note: BGP peers, NetBIOS, and BOOTP cannot be configured directly on an unnumbered interface.

For information on using Site Manager to configure a BGP peer-to-peer session on routers connected through unnumbered interfaces, see “Configuring BGP Peers over an Unnumbered Point-to-Point Link” on page 5-5.

To route NetBIOS packets over an unnumbered interface, you must configure a static entry to the Name Server.

To run BOOTP over unnumbered interfaces, you must select a preferred BOOTP server. For instructions, see *Configuring SNMP, BOOTP, DHCP, and RARP Services*.

Note: Unnumbered interfaces cannot be pinged directly. For this reason, such interfaces can add to the difficulty of diagnosing router problems.

As it does with routes learned over numbered interfaces, IP stores each route learned over an unnumbered interface in the routing table.

The routing-table entry for a route learned over an unnumbered interface contains the following values:

Next-hop address	0
Next-hop mask equals	0
Next-hop interface	Circuit number of the unnumbered interface

Using the Alternate Associated Address Option

The alternate associated address option ensures that a network on an unnumbered interface remains reachable. IP automatically assigns an alternate associated address to an unnumbered interface in the event that the primary associated address has gone down. IP uses the first available interface.

Note: In the event that an unnumbered associated address becomes unreachable, some functionality may be lost for certain protocols over the unnumbered interface.

Note: In some configurations, changing the associated address can affect the way routes are advertised. In a RIP configuration, for example, if you change the associated address for an unnumbered interface configured with RIP, you may change the way RIP advertises subnets.

You can use Site Manager to enable the associated address option.

Site Manager: UnNumbered Assoc Alternate parameter: page A-36

Chapter 3
Customizing IP Services

This chapter consists of the following sections:

- Configuring IP on the Router -- page 3-3

  Configuring the Router for Host-Only Mode -- page 3-4
  Setting the Time-to-Live Value on a Source Packet -- page 3-8
  Setting the RIP Diameter -- page 3-8
  Allowing an All-Zero or All-One Subnet Address -- page 3-9
  Estimating the Size of the Routing Table -- page 3-9
  Using a Default Route for an Unknown Subnet -- page 3-10
  Specifying the Maximum Number of IP Policies -- page 3-10
  Disabling Route Filter Support -- page 3-10
  Enabling Equal-Cost Multipath Support -- page 3-11
  Enabling ISP Mode on the Router -- page 3-12

- Configuring an IP Interface -- page 3-13

  Configuring a Multinet Interface -- page 3-15
  Specifying a Broadcast Address for an Interface -- page 3-16
  Specifying the Cost of an Interface -- page 3-17
  Enabling MTU Discovery on an Interface -- page 3-18
  Enabling ICMP Address-Mask Replies on an Interface -- page 3-19
  Enabling All-Subnet Broadcasting on an Interface -- page 3-20
  Configuring Address Resolution -- page 3-27
  Defining an Adjacent Host -- page 3-33
  Specifying a MAC Address or E.164 Address -- page 3-21
  Enabling Source Routing over a Token Ring Network -- page 3-22
  Enabling ICMP Redirect Messages -- page 3-19
  Configuring an SMDS Address -- page 3-24
  Configuring a WAN Address for a Frame Relay Network -- page 3-24
  Specifying the Maximum Size of the Forwarding Table -- page 3-25
  Configuring an Interface for an ATM Logical IP Subnet -- page 3-26
  Configuring an Interface for an ATM Logical IP Subnet -- page 3-26

- Configuring Address Resolution -- page 3-27

  Selecting an Address Resolution Scheme for an Interface -- page 3-29
  Selecting an Encapsulation Option for ARP and Probe -- page 3-30
  Enabling Proxy ARP on an Interface -- page 3-31
  Timing Out Entries in the Address Resolution Cache -- page 3-32
  Defining an Adjacent Host -- page 3-33

- Configuring IP Accounting on a Frame Relay Interface -- page 3-35

  Enabling IP Accounting on the Router -- page 3-35
  Specifying the Maximum Size of the IP Accounting Table -- page 3-35
  Controlling Notification of a Full IP Accounting Table -- page 3-36
  Copying the IP Accounting Table to the Checkpoint Table -- page 3-36

- Configuring the Trivial File Transfer Protocol -- page 3-36
- Configuring the Revised IP Security Option on an Interface -- page 3-37
- Defining a Static Route -- page 3-45
- Defining a Black Hole for a Supernet -- page 3-46
- Configuring Router Discovery -- page 3-46
- Connecting the Router to a Blacker Front End -- page 3-47

Configuring IP on the Router

When you add IP to the router, IP is automatically enabled with a set of default characteristics and services.

You can accept the default characteristics or use the IP global parameters described in the following sections to customize IP.

IP runs on all slots that contain IP interfaces.When you add an IP interface to a circuit, Site Manager enables IP on the slot. When you set IP global parameters, you are affecting the way IP operates on all slots where it is configured.

Use Site Manager to disable and re-enable IP on the router.

Site Manager: Enable parameter: page A-39

Configuring the Router for Host-Only Mode

By default, IP forwards all packets that are not addressed to itself. You can also configure IP in nonforwarding -- or *host-only* -- mode.

Use the forwarding mode if you want the IP router to route (forward) IP traffic. Forwarding configures the IP router to process all broadcast packets and all IP packets explicitly addressed to it, and to route all other IP packets.

Select nonforwarding mode on the router if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces but also want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access. Nonforwarding mode configures the IP router to act as an IP host; it does not forward IP traffic, but it still processes packets explicitly addressed to it. In Not Forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated.

Because the IP router does not forward IP traffic in Not Forwarding mode, you must configure the router to bridge IP traffic not explicitly addressed to it. You must configure the bridge for each circuit that conveys IP datagrams. The bridge will then forward all IP datagrams that are not explicitly addressed to the router.

You can use Site Manager to select nonforwarding (host-only) mode.

Site Manager: Forwarding parameter: page A-39

Table 3-1 and Table 3-2 show valid and invalid configurations for source route bridges and learning bridges. Each configuration is expressed as

source device > *medium* > *destination device*

For example:

Bay > Eth > Bay

source device is a Bay Networks router, the *medium* is Ethernet, and the *destination device* is a Bay Networks router.

A more complex configuration is expressed as

source device > *medium* > *intermediate device* > *medium* > *destination device*

For example:

Bay > Token > Bay > Token > ES

source device is a Bay Networks router, the first *medium* is a Token Ring network, the *intermediate device* is a Bay Networks router, the second *medium* is another Token Ring network, and the *destination device* is an endstation from another vendor.

Table 3-1. Source Route Bridge Support for Host-Only Mode

Bridge Configuration	Result
Bay* > Eth† > Bay	OK
Bay > Token‡ > Bay	OK
Bay > FDDI** > Bay	OK
Bay > PTP†† > Bay	OK
Bay > FR‡‡ > Bay	OK
Bay > SMDS*** > Bay	OK
Bay > PPP††† > Bay	Fails
Bay > Eth > Bay > Token > ES‡‡‡	Fails
Bay > Token > Bay > Token > ES	OK
Bay > FDDI > Bay > Token > ES	Fails
Bay > PTP > Bay > Token > ES	Fails
Bay > FR > Bay > Token > ES	Fails
Bay > SDMS > Bay > Token > ES	Fails
Bay > PPP > Bay > Token > ES	Fails
ES > Token > Bay > Eth > Bay	Fails
ES > Token > Bay > Token > Bay	OK
ES > Token > Bay > FDDI > Bay	Fails
ES > Token > Bay > PTP > Bay	Fails
ES > Token > Bay > FR > Bay	Fails

(continued)

Table 3-1. Source Route Bridge Support for Host-Only Mode *(continued)*

Bridge Configuration	Result
ES > Token > Bay > SDMS > Bay	Fails
ES > Token > Bay > PPP > Bay	Fails
ES > Token > Bay > Eth > Bay > Token > ES	OK
ES > Token > Bay > Token > Bay > Token > ES	OK
ES > Token > Bay > FDDI > Bay > Token > ES	OK
ES > Token > Bay > PTP > Bay > Token > ES	OK
ES > Token > Bay > FR > Bay > Token > ES	OK
ES > Token > Bay > SDMS > Bay > Token > ES	OK
ES > Token > Bay > PPP > Bay > Token > ES	OK

*. Bay Networks router with bridge and IP in host-only mode

†. Ethernet connection

‡. Token Ring connection

**. FDDI connection

††. Bay Networks proprietary point-to-point synchronous connection

‡‡. Frame Relay synchronous connection

***. SMDS synchronous connection

†††. PPP synchronous connection

‡‡‡. Station you are communicating to or from if not Bay Networks

Table 3-2. Learning Bridge Support for Host-Only Mode

Bridge Configuration	Result
Bay* > Eth† > Bay	OK
Bay > Token‡ > Bay	OK
Bay > FDDI** > Bay	OK
Bay > PTP†† > Bay	OK
Bay > FR‡‡ > Bay	OK
Bay > SMDS*** > Bay	OK
Bay > PPP††† > Bay	Fails
Bay > Eth > Bay > Eth > ES‡‡‡	OK
Bay > Token > Bay > Eth > ES	Fails

(continued)

Table 3-2. Learning Bridge Support for Host-Only Mode *(continued)*

Bridge Configuration	Result
Bay > FDDI > Bay > Eth > ES	OK
Bay > PTP > Bay > Eth > ES	OK
Bay > FR > Bay > Eth > ES	Fails
Bay > SDMS > Bay > Eth > ES	Fails
Bay > PPP > Bay > Eth > ES	Fails
ES > Eth > Bay > Eth > Bay	OK
ES > Eth > Bay > Token > Bay	Fails
ES > Eth > Bay > FDDI > Bay	OK
ES > Eth > Bay > PTP > Bay	OK
ES > Eth > Bay > FR > Bay	Fails
ES > Eth > Bay > SDMS > Bay	Fails
ES > Eth > Bay > PPP > Bay	Fails
ES > Eth > Bay > Eth > Bay > Eth > ES	OK
ES > Eth > Bay > Token > Bay > Eth > ES	OK
ES > Eth > Bay > FDDI > Bay > Eth > ES	OK
ES > Eth > Bay > PTP > Bay > Eth > ES	OK
ES > Eth > Bay > FR > Bay > Eth > ES	OK
ES > Eth > Bay > SDMS > Bay > Eth > ES	OK
ES > Eth > Bay > PPP > Bay > Eth > ES	OK

*. Bay Networks router with bridge and IP in host-only mode

†. Ethernet connection

‡. Token Ring connection

**. FDDI connection

††. Bay Networks proprietary point-to-point synchronous connection

‡‡. Frame Relay synchronous connection

***. SMDS synchronous connection

†††. PPP synchronous connection

‡‡‡. Station you are communicating to or from if not Bay Networks

Setting the Time-to-Live Value on a Source Packet

Each IP data packet includes a *time-to-live* (TTL) value. The TTL value specifies the maximum number of hops that the packet is allowed to traverse in the network before an intermediate router discards it.

The router that originates the packet sets the TTL to a positive value. Each router that receives the packet decrements the TTL counter by one. A router that receives a packet with a TTL of zero discards the packet -- if the packet is not addressed to itself. The TTL counter prevents packets from looping endlessly through the network.

By default, IP sets the TTL field on each source packet (that is, each packet that it originates and transmits) to 30 hops. You can use Site Manager to specify a TTL value (overriding the default).

Site Manager: Default TTL parameter: page A-41

Setting the RIP Diameter

The *RIP diameter* is a hop count that the Routing Information Protocol (RIP) uses to denote infinity. In order for RIP to operate properly, every router within the network must be configured with an identical RIP diameter value. If RIP is enabled, this parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, the IP router still uses the RIP diameter to determine network width.

The default RIP diameter value is 15 hops. You must set this parameter so that none of the interface cost, static cost, or route filter cost parameters exceed the RIP diameter. We recommend that you accept the default RIP diameter value.

Use Site Manager to specify the RIP diameter.

Site Manager: RIP Diameter parameter: page A-42

Allowing an All-Zero or All-One Subnet Address

By default, IP considers an interface with an address whose subnet portion consists of all zeros or all ones to be an illegal address.

Using Site Manager, you can configure IP to allow an all-zero and all-one subnet address.

Site Manager: Zero Subnet Enable parameter: page A-42

Enable this feature with caution: it can result in an ambiguous address. For example, if an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all-subnets broadcast from a directed broadcast for the zero subnet.

Estimating the Size of the Routing Table

IP allows you to estimate the number of networks and hosts that will require an entry in the IP routing table. The router uses your estimate to preallocate memory for the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation. Preallocation also makes better use of memory and reduces the amount of memory required. By default, the router allocates resources to support 500 network and host entries in the routing table.

If you have enabled ISP mode, the default value is 40,000 entries. You must reduce this value to an appropriate size if the sysem is running with 8 MB or 16 MB processor modules. Failure to change the value will result in an out-of-memory error on these processors.

You can use Site Manager to estimate the number of networks and hosts (overriding the default).

Site Manager: Estimated Networks parameter: page A-43
Site Manager: Estimated Hosts parameter: page A-43

Avoid making an estimate that is excessively large. This will cause a wasteful overallocation of memory.

Using a Default Route for an Unknown Subnet

By default, IP drops packets addressed to an unknown subnet and returns an ICMP to the sender. This prevents local traffic from accidentally following the default route to the Internet.

In cases where remote sites follow a default route to a central site, it is appropriate to enable this parameter. If the router serves as an Internet gateway (with a default route to the Internet), the parameter can be disabled.

You can use Site Manager to configure IP to use a default route for packets addressed to an unknown subnet.

Site Manager: Enable Default Route for Subnets parameter: page A-44

The default route must be present in the routing table.

Specifying the Maximum Number of IP Policies

By default, IP allows you configure up to 32 announce policies and 32 accept policies for each protocol that you configure.on the router. This value must be increased for IP to allocate memory and implement additional policies.

You can use Site Manager to configure more than 32 Accept or Announce policy rules for a protocol.

Site Manager: Maximum Policy Rules parameter: page A-44

Disabling Route Filter Support

By default, IP supports route filters. You can use Site Manager to disable route filter support When route filter support is disabled, IP does not allocate memory for route filters when the maximum number of IP policies is increased..

Site Manager: Route Filter Support parameter: page A-44

Enabling Equal-Cost Multipath Support

IP equal-cost multipath support is a load-balancing feature that allows IP to distribute traffic over multiple (up to five) equal-cost paths to the same destination.

By default, multipath support is disabled on the router: for each routing protocol, IP stores the best next hop to a destination in the routing table. If traffic arrives on an interface, IP determines the best route to the destination and forwards all packets out the next-hop interface.

If multipath support is enabled, IP will distribute traffic among equal-cost routes (if such routes are available). You can use Site Manager to enable multipath support on the router and to specify whether IP distributes packets in a round-robin fashion or uses a distribution method based on the source and destination address of the packets.

- In *round-robin distribution*, IP forwards each packet to a different next hop until it reaches the end of the list of available next hops; then it repeats the list. Round robin distribution makes full use of available resources but may cause packets to be delivered out of order.
- In *address-based distribution* based on the source and destination address, IP forwards all packets with a given source and destination address to the same next hop. This method increases the chances that the packets will be delivered in order but does not take full advantage of available resources.
- In address-based distribution based on the destination address only, IP forwards all packets with a given destination address to the same next hop. This method makes the forwarding algorithm compatible with RSVP when RSVP is enabled on the router.

Site Manager: Multiple Next Hop Calculation Method parameter: page A-45

Both RIP and OSPF support equal-cost multipath routing. You can use Site Manager to specify the number of equal-cost routes (up to five) that RIP and OSPF will submit to the routing table manager. For instructions, see “Configuring RIP for Equal-Cost Multipath Support” on page 4-8 and “Configuring OSPF for Equal-Cost Multipath Routing” on page 5-13.

Enabling ISP Mode on the Router

IP provides certain features designed for use by internet service providers (ISPs). This set of features optimizes the router forwarding performance in a core Internet backbone application and includes a BGP soloist.

By default ISP mode features are disabled on the router. You can use Site Manager to enable the following features:

- *Suppression of IP forwarding tables*. By default, IP maintains a forwarding table on each IP interface. IP maintains this table as a cache for routes that are frequently used to forward data packets that arrive on the interface. However, if the number of frequently-used routes exceeds the size of the forwarding table, the router will be continually updating the forwarding table by removing old routes and installing new route entries. ISP mode disables all forwarding tables on all interfaces and optimizes the routing table to allow direct forwarding, avoiding the overhead of cache misses and cache updates. If you select ISP mode, you do not have to explicitly disable the forwarding tables on each interface.
- *BGP soloist.* By default, BGP runs on all slots configured with IP interfaces. In ISP mode, BGP runs as a soloist. (For further information see “Configuring BGP as a Soloist” on page 6-13.)

Site Manager: ISP Mode parameter: page A-45

The following parameter settings also help optimize the router’s performance and operation:

IP Global Parameter	Setting
Route Filter Support	Disabled
Maximum Policy Rules	Set ss required
Estimated Networks	Set as required
ICMP Redirects	Set to off at router interconnection points

Configuring an IP Interface

An IP network *interface* consists of a physical circuit configured with the appropriate data link and IP protocols. Each interface connects the router to one or more IP networks.

For example, the router in Figure 3-1 is configured with three IP interfaces. One of these interfaces is a point-to-point interface that connects the router to a single long-haul medium terminated by a host or another router. The other two interfaces are LAN interfaces that connect the router to an Ethernet or FDDI local area medium.

An IP interface can provide access to multiple networks. For example, in Figure 3-1, LAN interface 1 provides a connection to both LAN B and LAN C.

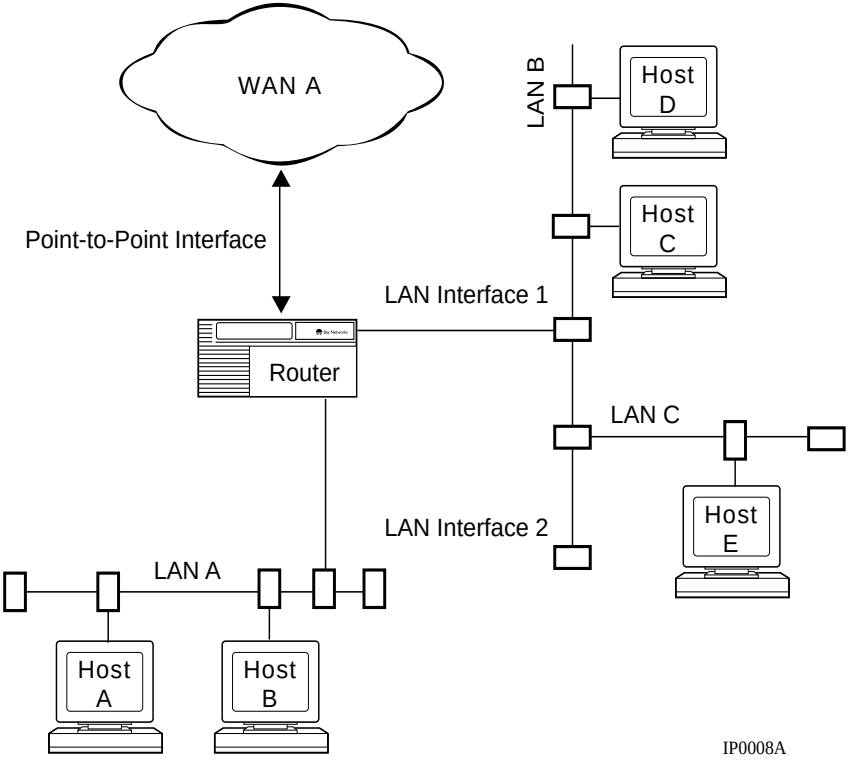

Figure 3-1. IP Interface

You associate a network with an interface by assigning the network's unique IP address to the circuit on which the interface is configured.

You can use Site Manager to assign an IP address and subnet mask to an interface.

Site Manager: IP Address parameter: page A-22
Site Manager: Subnet Mask parameter: page A-22

When you add IP to an interface, IP is automatically enabled on that interface. You can use Site Manager to disable IP on the interface.

Site Manager: Enable parameter: page A-24

Note: When you reconfigure an interface in dynamic mode, IP restarts on that interface. Thus, if the interface you reconfigure is the interface that supports Site Manager's SNMP connection to the router, restarting IP on that interface will cause Site Manager to temporarily lose its router connection and to display a warning message. To verify that the change took effect, display the IP Global Parameters window and inspect the setting.

If you are configuring IP over an SMDS circuit, be sure to enter the correct addresses in the MAC Address, SMDS Group Address, and SMDS Arp Req Address parameter boxes displayed on this screen. These addresses are the same as those you entered in the Individual Address, Group Address, and ARP Address parameters of the SMDS Configuration window when you configured SMDS.

Configuring a Multinet Interface

The *multinet* capability allows you to assign multiple IP network/subnet addresses to a single circuit; each IP address represents a separate network interface on the circuit.

Multinet is commonly used in IP networks as part of a transition strategy. As networks evolve it is sometimes necessary to consolidate several physical networks. To avoid renumbering, the physical networks can be consolidated onto a multinetted router interface. This allows hosts to migrate to the new IP interface or maintain the old IP address.

In Figure 3-2, for example, Host A and Host C are located on different subnets. The router will provide connectivity between Hosts A and C by acting as the default gateway and routing packets.

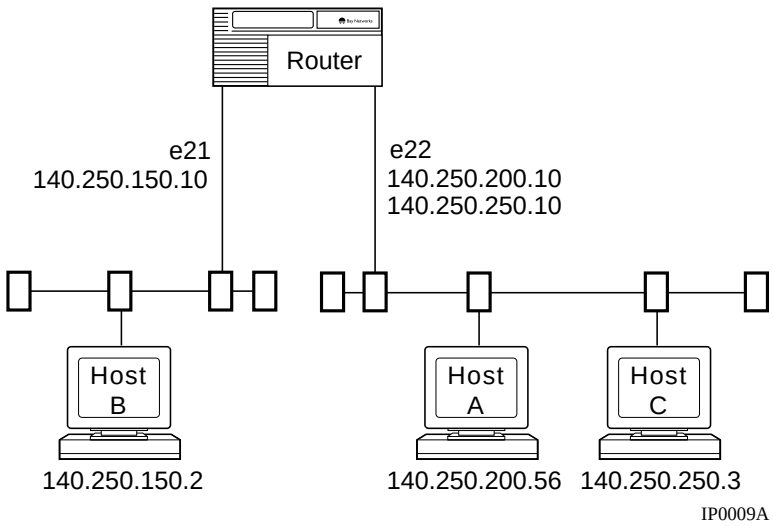

Figure 3-2. Multinet Configuration

Specifying a Broadcast Address for an Interface

In broadcasting, the IP router transmits a single packet to every host on an attached network. To do so, it uses a broadcast address that refers to all hosts on the network. A broadcast address is simply an IP address that contains all 1s or all 0s in the host portion.

For example, the IP Class C address 10.3.45.12 has the following characteristics:

- Because the address is for a Class A network (the network portion is 1 byte), the host portion contains 3 bytes.
- Because the host portion of a broadcast address consists of all 1s or all 0s, the broadcast address for that network can be one of the following: 10.255.255.255, 10.0.0.0, 255.255.255.255, or 0.0.0.0.

Some networks do not support broadcasts; thus, configuring an IP broadcast address does not guarantee efficient broadcast delivery.

By default, IP uses a broadcast address that contains all 1s in the host portion.

You can use Site Manager to configure a broadcast address on an IP interface.

Site Manager: Broadcast Address parameter: page A-25.

Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted decimal notation.

Specifying a Subnet Broadcast Address

The way you configure a broadcast address for a subnet is different from the way you configure a broadcast address for a network. When you extend the network portion of the IP address to create a subnet address, you automatically take away from the host portion of the address. To configure a subnet broadcast, you take the subnet mask for that subnet and invert it. For example, if the IP address of the subnet is 10.4.2.3, and the mask is 255.255.0.0, then the subnet broadcast address is either 10.4.255.255 or 10.4.0.0.

Specifying the Cost of an Interface

Each IP interface has an assigned cost. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets transmitted out other interfaces.

By default, an IP interface has a cost of 1.

You can use Site Manager to specify a cost (overriding the default).

Site Manager: Interface Cost parameter: page A-25

Enter the interface cost value (standard RIP implementation assigns a cost of 1); however, keep in mind that increasing this value causes the upper bound set by the RIP Network Diameter parameter to be attained more rapidly.

Enabling MTU Discovery on an Interface

A Probe MTU is a request for the maximum transmission unit (MTU) size used on all networks an IP datagram must traverse from source to destination.

By default, IP does not respond to Probe Requests. Using Site Manager, you can configure IP to issue a Reply MTU on this interface in response to a Probe MTU.

Site Manager: MTU Discovery parameter: page A-26

By enabling IP to respond to Probe MTUs on this interface, you eliminate transit fragmentation and destination reassembly for datagrams destined for this interface and, therefore, decrease network load.

The Reply MTU and the Probe MTU are options 11 and 12 in RFC 1063.

Configuring ICMP Features on an Interface

IP routers and hosts exchange control messages defined by the Internet Control Message Protocol (ICMP).

The following sections describe the ICMP messages that you can enable and disable on an IP interface.

Enabling ICMP Address-Mask Replies on an Interface

You can configure IP to generate ICMP (Internet Control Message Protocol) address-mask-reply messages on this interface in response to valid address-mask-request messages. The interface generates ICMP address-mask-reply messages in compliance with the relevant sections of RFCs 950 and 1009.

By default, IP does not generate address-mask-reply messages. You can use Site Manager to enable the feature on the interface.

Site Manager: Addr Mask Reply parameter: page A-26

Enabling ICMP Redirect Messages

An ICMP redirect is a message sent by the router to alert a host that it should be using a different path to route data.

By default, IP sends ICMP redirect messages. You can use Site Manager to disable and re-enable this function.

Site Manager: Redirects parameter: page A-30

In some cases, you do not want an interface to send out redirects. For example, in a Frame Relay network, two stations on the same network may not be directly connected if the network is not fully meshed. Thus, in this case, you would set Redirects to Disable.

Enabling All-Subnet Broadcasting on an Interface

An all-subnet broadcast (ASB) datagram has a destination address equal to the broadcast address for an entire network (all subnets). For example, if a network interface serves the subnet 128.10.2.1 with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram.

By default, IP does not flood ASB datagrams. Using Site Manager, you can configure IP to flood ASB datagrams out this interface:

Site Manager: All Subnet Bcast parameter: page A-27

Disabling UDP Checksum Processing on the Interface

By default, UDP checksum processing is enabled on this interface. All outgoing and incoming UDP datagrams are subject to checksumming.

You can use Site Manager to disable UDP checksum processing.

Site Manager: Checksum parameter: page A-29

You should use the default in virtually all instances. Disable UDP checksum processing to provide backward compatibility with UNIX BSD 4.1.

Specifying a MAC Address or E.164 Address

You can use Site Manager to specify a MAC address or an E.164 address for this interface.

Site Manager: MAC Address parameter: page A-29

0	The IP router uses its IP address and the circuit's MAC address when transmitting packets on this interface.
A user-specified MAC address	The IP router uses its IP address and this MAC address when transmitting and receiving packets on this interface.
An E.164 address.	If the interface is on an SMDS circuit, by default, IP uses the individual SMDS-configured address. You can enter the entire E.164 address—for example, C1 617 555 5000 FFFF.

To configure this parameter for a multinet or multigroup configuration, refer to *Configuring SMDS*.

Enabling Source Routing over a Token Ring Network

The IP router can route over Token Ring (TR) networks that contain one or more source-routing bridges.

In a source-routing network, every endstation that sends out a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for IP routers to route packets across a source-routing network, *they must act like endstations*, supplying route descriptors for each packet before they send it out onto the network.

With endnode support enabled, whenever an IP router receives a packet and determines that the packet's next hop is located across a source-routing network, the router

- Adds the necessary Routing Information Field (RIF) information to the packet's MAC header
- Sends the packet out onto the network where it is source routed toward the next hop

Upon receiving the packet from the Token Ring network, the peer router strips off the RIF and continues to route the packet toward the destination network address (Figure 3-3).

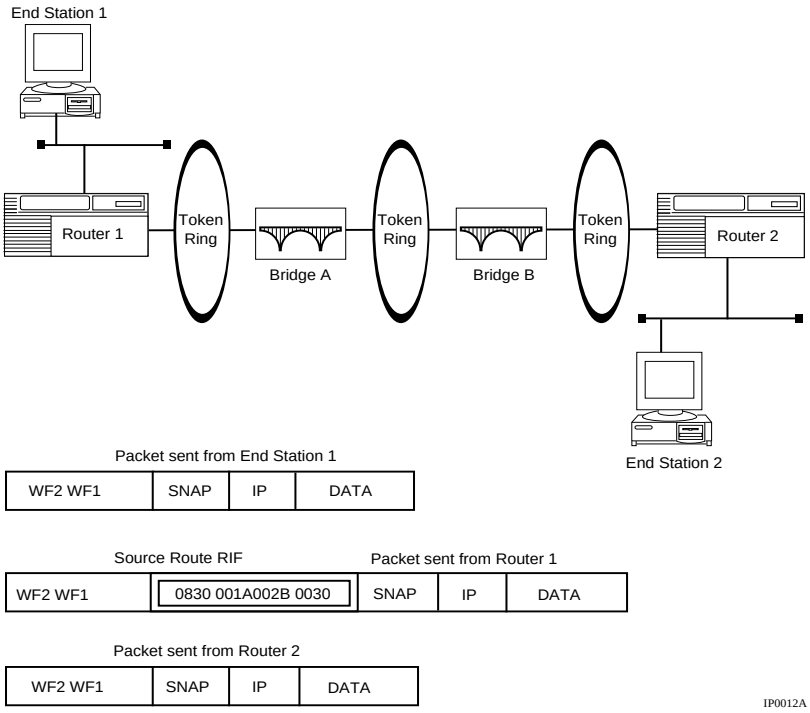

Figure 3-3. IP Routers Source Routing across a Token Ring Network

The router can send ARP packets over an interface configured for a Token Ring network. Bay Networks supports both spanning tree explorer (STE) and all route explorer (ARE) ARP packets.

You can use Site Manager to configure source-route endnode support on a per-circuit basis. Site Manager also allows you to select STE or ARE ARP packets.

Site Manager: TR Endstation parameter: page A-30.
Site Manager: TR Endstation ARP Type parameter: page A-38.

Configuring an SMDS Address

By default, if the interface is connected to an SMDS network, IP uses the SMDS-configured addresses.

You can use Site Manager to supply

- A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider
- An address-resolution multicast address for this IP interface in an SMDS network

Site Manager: SMDS Group Address parameter: page A-31
Site Manager: SMDS Arp Req Address parameter: page A-32

Configuring a WAN Address for a Frame Relay Network

If the interface is connected to a Frame Relay Network, you can use Site Manager to configure

- A broadcast address. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will broadcast the message.
- A multicast address for this IP interface that will send messages to all OSPF routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF routers.
- A multicast address for this IP interface that will send messages to all OSPF designated routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF designated routers.

Site Manager: WAN Broadcast parameter: page A-32
Site Manager: WAN Multicast #1 parameter: page A-33
Site Manager: WAN Multicast #2 parameter: page A-33

Specifying the Maximum Size of the Forwarding Table

To keep to a minimum the amount of time it spends looking up routes, IP creates and maintains a cache of frequently used routes -- called a *forwarding table* -- for each IP interface.

A forwarding table is a first-in-first-out (FIFO) buffer. When a datagram arrives on an IP interface for forwarding, IP performs the following steps:

1. IP searches the forwarding table associated with the interface for the destination network.
2. If the search is successful, IP dispatches the datagram to the interface noted in the table entry.
3. If the search is unsuccessful, IP consults the routing table to get the same information, dispatches the datagram to the appropriate interface, and caches the information in the appropriate forwarding table -- either by appending information to the table (if the table is not full) or by overwriting the oldest, first-in table entry (if the table is full).

If IP flushes a route from the routing table, it also removes the route from the forwarding tables, thus ensuring that invalid routing information is not retained in interface-specific caches.

An interface that receives packets that are destined for a large number of different destinations may benefit from a larger forwarding table. The larger the number of entries, the more likely it is that the destination will already be in the forwarding table and the faster the route lookups will be for those destinations.

Keep in mind that configuring a forwarding table size that is larger than necessary reduces the total amount of memory usable by other applications. Configuring a routing table too small can affect overall router performance. A check of the number of cache hits and misses will help determine the optimal size of the forwarding table. For debugging purposes, if you see the wfIpInterfaceCacheMisses statistic going up at a rapid rate, you should consider increasing the table size. However, an occasional cache miss does not warrant an increase in table size.

You can use Site Manager to control the number of destinations that are cached in the forwarding table on this receiving interface.

Site Manager: Max Forwarding Table Size parameter: page A-35

Configuring an Interface for an ATM Logical IP Subnet

RFC 1577, “Classical IP and ARP over ATM.” is a specification for an administrative entity within an ATM network called a logical IP subnet (LIS). Bay Networks supports RFC 1577. For information on configuring IP interfaces on an ATM LIS, see *Configuring ATM Services.*

Using Site Manager, you can

- Specify whether the router is running as an ATM client or server on this interface. You must configure one ATMARP server for each logical IP subnet you define.
- Define the ATM address network prefix of the ATMARP server on your network. A complete ATM address consists of a network prefix and a user part.
- Define the user part (suffix) of the ATM address for the ATMARP server on your network. The user part suffix consists of a 6-byte endstation identifier and a 1-byte selector field.
- Specify (for a client) the interval between registration refreshes.
- Specify (for a server) the duration for which the registration is valid.

Site Manager: ATM ARP Mode parameter: page A-37
Site Manager: ARP Server ATM Address Network Prefix parameter: page A-37
Site Manager: ARP Server ATM Address User Part parameter: page A-38
Site Manager: Registration Refresh Interval parameter: page A-38

Configuring Address Resolution

The IP router needs both a physical address and an IP address to transmit a datagram. In situations where the router knows only the network host's IP address, the Address Resolution Protocol (ARP) enables the router to determine a network host's physical address by binding a 32-bit IP address to a 48-bit MAC address. A router can use ARP across a single network only, and the network hardware must support physical broadcasts.

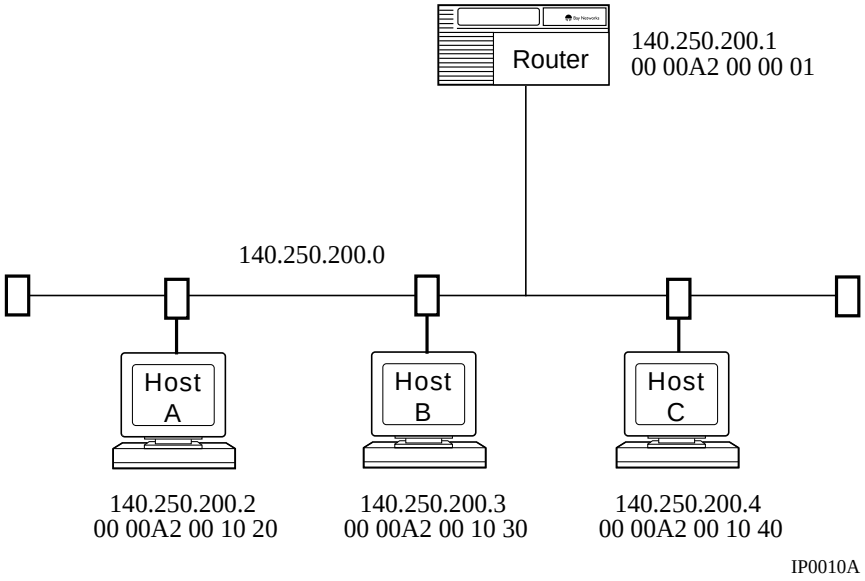

Figure 3-4. ARP Example

For example, in Figure 3-4, the router and Host C are on the same physical network. Both devices have an assigned IP address (the router's is 140.250.200.1 and Host C's is 140.250.200.4) and both devices have an assigned physical address (the router's is 00 00 A2 00 00 01 and Host C's is 00 00 A2 00 10 40).

In Figure 3-4, the router wants to send a packet to Host C but knows only Host C's IP address. The router uses ARP to determine Host C's physical address, as follows:

1. The router broadcasts a special packet, called an ARP request, that asks IP address 140.250.200.4 to respond with its physical address.
2. All network hosts receive the broadcast request.
3. Only Host C responds with its hardware address.

The router maps Host C's IP address (140.250.200.4) to its physical address (00 00 A2 00 10 40) and saves the results in an address-resolution cache for future use.

Note: It is possible for the router to send out ARP requests even if ARP, which is a dynamically loaded module, is not currently loaded on the router. It is the responsibility of the network administrator to ensure that ARP is loaded correctly on a slot. To do this through Site Manager, select Events Manager > Options > Filters; then select LOADER and Debug, and do a File > Get Current Log File. Verify that ARP is loaded on a slot by locating the following message in the log:

```
#xx:01/01/95 10:10:55.00 DEBUG SLOT x LOADER CODE:33
Loader service completed for ARP.EXE 0xxxxxxxxx
```

You can use Site Manager to

- Control how ARP acts in relation to IP's forwarding state.
- Control whether IP drops and logs an invalid ARP source address or simply drops the request.
- Control whether IP drops or accepts ARP requests in which the source and destination addresses are located in different networks or subnetworks. This parameter allows Proxy ARP to generate replies when the source and destination networks in the ARP request differ.

Site Manager: ARP Forwarding parameter: page A-40
Site Manager: Nonlocal ARP Source parameter: page A-40
Site Manager: Nonlocal ARP Destination parameter: page A-41

Selecting an Address Resolution Scheme for an Interface

In addition to ARP, IP supports Inverse ARP, HP Probe, and X.25 address resolution schemes:

- Inverse ARP provides address resolution for Frame Relay interfaces. It is used to discover the IP address of the station at the remote end of the virtual circuit.
- HP Probe, a Hewlett-Packard proprietary protocol, is an address resolution mechanism that functions much like ARP to determine a network host's physical address when all it knows is the network host's IP address, by binding a 32-bit IP address to a 48-bit MAC address. IP supports HP Probe over Ethernet and the following HP Probe messages:

  Unsolicited Reply (incoming and outgoing)
  Name Request (incoming)
  Name Reply (outgoing)
  Virtual Address Reply (incoming and outgoing)
  Virtual Address Request (incoming and outgoing)
  Proxy Request (incoming and outgoing)
  Proxy Reply (incoming and outgoing)

Note: If bridging is configured and enabled on the interface (in addition to IP), the Name Request/Reply and the Proxy Request/Reply messages are bridged.

  IP can support the concurrent operation of HP Probe and ARP on an interface.
- The X.25 address resolution scheme is used on network interfaces that support the X.25 DDN service.
- The RFC 877-compliant address-resolution mechanism is used on network interfaces that support the X.25 PDN service.

On interfaces configured for a Token Ring network, the router can send ARP requests as Spanning Tree Explorer (STE) packets or All-Routes Explorer (ARE) packets.

By default, ARP is enabled on the interface. Use Site Manager to specify any of the following supported address resolution schemes.

Site Manager: Address Resolution parameter: page A-27

Selecting an Encapsulation Option for ARP and Probe

If you select ARP, Probe, or ARP/Probe, you must also select the appropriate datalink encapsulation option.

- If your address-resolution scheme is ARP only, select Ethernet encapsulation, SNAP encapsulation, or Ethernet/SNAP encapsulation.
- If your resolution scheme is HP Probe only, select LSAP encapsulation.
- If your resolution scheme is ARP/Probe, select Ethernet/LSAP encapsulation, SNAP/LSAP encapsulation, or Ethernet/SNAP/LSAP encapsulation.

By default, IP uses ARP Ethernet encapsulation. You can use Site Manager to specify an encapsulation scheme.

IP ignores this parameter if the underlying medium is anything other than Ethernet.

Site Manager: Enet Arp Encaps parameter: page A-31

Enabling Proxy ARP on an Interface

Proxy ARP allows a router to answer a local ARP request for a remote destination. For example, in Figure 3-5, Hosts B and C are located on the same network but on separate subnetworks. Hosts B and C do not understand subnetworking. The router connecting the two physical networks knows which host resides on which network. The address mask is 255.255.255.000. In this example, one subnet is a remote network with respect to the other subnet.

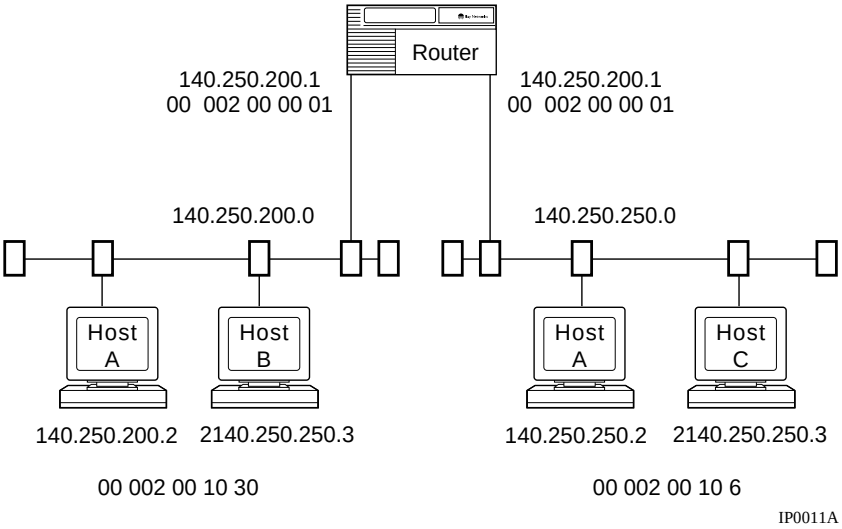

Figure 3-5. Proxy ARP Example

Host B wants to talk to Host C, so Host B broadcasts an ARP request, which asks IP address 140.250.250.2 to respond with its physical address. The router captures Host B's ARP request and responds with its hardware address 00 00 A2 00 00 01 and Host C's IP address 140.250.250.2. Host B maps Host C's IP address 140.250.250.2 to the router's hardware address 00 00 A2 00 00 01.

By default, Proxy ARP is disabled on the interface. You can use Site Manager to enable Proxy ARP.

Site Manager: Proxy parameter: page A-28

With Proxy ARP enabled, the router will respond with an ARP reply if there is a valid route (that is, if the router is able to forward traffic) to the destination in the routing table. This route may be a subnet route or a default route. For the router to respond for subnets that are reachable via the default route, you must use Site Manager to configure IP to use a default route for unknown subnets.

Site Manager: Enable Default Route for Subnets parameter: page A-44

There are devices that use Proxy ARP to determine a gateway rather than relying on a statically defined default gateway. These devices will use ARP for all remote destinations.

To enable the router to reply to ARP for remote destinations on other networks, you must enable Proxy ARP and use Site Manager to set the Nonlocal ARP Destination parameter to Accept.

Site Manager: Nonlocal ARP Destination parameter: page A-41

Timing Out Entries in the Address Resolution Cache

IP maintains an address resolution cache on each interface that is configured with ARP or Proxy Arp. The address resolution cache contains host physical addresses learned by means of ARP or Proxy ARP.

If you enable the cache timeout feature on this interface, the IP router removes address resolution cache entries that have not been accessed within the specified number of seconds. Once an entry is removed, the IP router must use ARP to re-acquire the physical-level address.

By default, the cache timeout feature is disabled on the interface. You can use Site Manager to enable the feature and to specify a timeout interval (in seconds).

Site Manager: Host Cache parameter: page A-28

A host entry is timed out (deleted) if the IP router sends no traffic destined for that host within the specified timeout period.

Defining an Adjacent Host

An *adjacent host* is a device on a locally attached network. This device may or may not be a router. You must configure a media access control (MAC) address for each adjacent host that does not implement the Address Resolution Protocol (ARP).

Also, if a local network does implement ARP, you may want to configure a MAC address for an adjacent host to pre-empt the ARP process.

You can use Site Manager to specify and supply

- The state (active or inactive) of the adjacent host in the IP routing tables.
- The IP address of the device for which you want to configure an adjacent host.
- The IP address of the router's network interface to the adjacent host.
- The subnet mask of the IP address specified for the Next Hop Addr parameter.
- The physical address of the adjacent host. This value can be a 48-bit Ethernet address, a 64-bit SMDS address, an ATM PVC VPI/VCI address, or, for an ATM SVC, the address of the ATM interface.
- The adjacent host's encapsulation method.
- The X.121 address of the adjacent host (for a PDN/X.25, DDN/X.25, or BFE/X.25 connection).

Site Manager: Enable parameter: page A-51
Site Manager: IP Adjacent Host Address parameter: page A-51
Site Manager: Next Hop Interface Addr parameter: page A-51
Site Manager: Next Hop Interface Mask parameter: page A-52
Site Manager: MAC Address parameter: page A-52
Site Manager: Host Encapsulation parameter: page A-52
Site Manager: Adjacent Host X.121 Address parameter: page A-53

Configuring IP Accounting on a Frame Relay Interface

IP Accounting is a mechanism for counting transit data packets -- that is, packets that IP receives on one interface and forwards out another interface. This mechanism allows a network service provider to bill a network user according to the amount of data that it routes between two locations.

Bay Networks currently provides IP Accounting support for Frame Relay networks. IP Accounting counts all data packets that the router receives on any IP interface and forwards out an IP/Frame Relay interface.

IP accounting keeps track of transit data packets by making an entry for each packet in an IP accounting table. Each entry includes the following fields: the source address of the packet, the destination address of the packet, the number of packets forwarded, and the number of bytes forwarded.

Warning: If the Frame Relay interface becomes overrun and the driver drops packets, these packets will still be counted by IP Accounting.

Enabling IP Accounting on the Router

By default, IP Accounting support is disabled on the router. You can use Site Manager to enable IP Accounting.

Site Manager: Enable parameter: page A-46

Specifying the Maximum Size of the IP Accounting Table

By default, the IP Accounting table can contain up 512 entries per slot.

You can use Site Manager to specify the maximum number of entries in the IP Accounting table.

Site Manager: Threshold parameter: page A-46

Controlling Notification of a Full IP Accounting Table

By default, IP Accounting sends a log message when the active IP Accounting table is 80 percent full. You must configure a trap to be sent. Use Site Manager to configure a trap exception for Entity 6 and event 99.

You can use Site Manager to specify a value from 1 to 100 (indicating the percentage of the maximum size) that causes IP Accounting to send a trap message.

Site Manager: Trap Percent parameter: page A-46

Copying the IP Accounting Table to the Checkpoint Table

Once the IP Accounting table is filled to capacity, IP Accounting can make no further entries until you empty the table. You empty the accounting table by copying its contents to a checkpoint table.

Using Site Manager:

1. **Open the IP Global window.**
2. **Click on Flag CKPT.**

IP Accounting maintains two aging counters, one for the accounting table and one for the checkpoint table. When you copy the contents, IP Accounting resets and starts both counters at 0.

Configuring the Trivial File Transfer Protocol

The Trivial File Transfer Protocol (TFTP) is a TCP/IP standard protocol for transferring files with minimum capability and minimal overhead. TFTP is implemented on top of the unreliable connectionless datagram delivery service and is used to move files between network devices.

TFTP was designed to be small and easy to implement. Because it is small, it is more restrictive, lacking most of the features of the File Transfer Protocol (FTP). TFTP provides inexpensive, unsophisticated file-transfer service only. It cannot list directories and provides no authentication.

TFTP runs on top of the User Datagram Protocol (UDP) and uses timeout and retransmission to ensure that data arrives. Each file transfer begins with a request to read or write to a file; this request also serves to ask for a connection. If the server grants the request, the connection is opened and the file is sent in fixed-length blocks (data packets) of 512 bytes. Each data packet contains one block of data and must be acknowledged by an acknowledgment packet before the next packet is sent. A data packet of less than 512 bytes terminates the transfer.

If a packet gets lost in the network, the intended recipient will time out and may retransmit its last packet (which can be data or an acknowledgment), causing the sender of the lost packet to retransmit the packet. Because the lock-step acknowledgment guarantees that all older packets have been received, the sender keeps one packet only on hand for transmission.

Both devices involved in a TFTP transfer are senders and receivers. One device sends data and receives acknowledgments; the other device sends acknowledgments and receives data.

The IP router includes a client and server implementation of TFTP, enabling the router to transmit and receive files across an Internet.

You can use Site Manager to specify the operating characteristics of TFTP on the router.

Site Manager: Enable parameter: page A-53
Site Manager: Default Volume parameter: page A-53
Site Manager: Close Time Out parameter: page A-54
Site Manager: Retry Time Out parameter: page A-54
Site Manager: Retransmit parameter: page A-54

Configuring the Revised IP Security Option on an Interface

IP routers support the Department of Defense (DoD) Revised IP Security Option (RIPSO), as defined in RFC 1108, on a per-interface basis. While RIPSO RFC 1108 specifies both “basic” and “extended” security options, our implementation supports only the basic option.

RIPSO is a feature that allows end systems and intermediate systems (routers) to add labels to or process security labels in IP datagrams that they transmit or receive on an IP network. The labels specify security classifications (for example, Top Secret Confidential, and Unclassified, in descending order), which can be used to limit the devices that can access these labeled IP datagrams.

As a labeled IP datagram traverses an IP network, only those systems that have the proper clearance (that is, whose security classification range covers the classification specified by the datagram) should accept and forward the datagram.

Any system whose security classification range does not cover the classification specified by the security label should drop the datagram.

Note: RIPSO does not include any method of preventing a system that does not support RIPSO from simply accepting and forwarding labeled datagrams. Thus, in order for RIPSO to be effective, *all* systems in a network must support RIPSO and process IP datagrams as described.

By default, RIPSO is disabled on IP interfaces. You can use Site Manager to enable RIPSO on an IP interface and specify the following:

- A range of acceptable security levels for IP datagrams the interface receives and transmits
- A set of required and allowed authority values for IP datagrams the interface receives and transmits
- Whether inbound datagrams received on this interface require security labels
- Whether outbound datagrams transmitted on this interface (either forwarded or originated by the router) require security labels
- Whether datagrams received or transmitted on this interface should have their labels stripped

You also specify whether the router creates the following types of labels:

- An implicit label, which the router uses to label unlabeled inbound datagrams, when required
- A default label, which the router uses to label unlabeled outbound datagrams, when required
- An error label, which the router uses to label ICMP error messages associated with processing security options

Site Manager: Enable Security parameter: page A-55
Site Manager: Strip Security parameter: page A-55
Site Manager: Require Out Security parameter: page A-56
Site Manager: Require In Security parameter: page A-56
Site Manager: Min Level parameter: page A-57
Site Manager: Max Level parameter: page A-57
Site Manager: Must Out Authority parameter: page A-58
Site Manager: Must In Authority parameter: page A-59
Site Manager: May In Authorit parameter: page A-59
Site Manager: Implicit Label parameter: page A-60
Site Manager: Implicit Authority parameter: page A-60
Site Manager: Implicit Level parameter: page A-61
Site Manager: Default Label parameter: page A-61
Site Manager: Default Authority parameter: page A-62
Site Manager: Error Label parameter: page A-63
Site Manager: Error Authority parameter: page A-63

Security Label Format

A RIPSO security label is three or more bytes long and specifies the security classification level and protection authority values for the datagram (Figure 3-6).

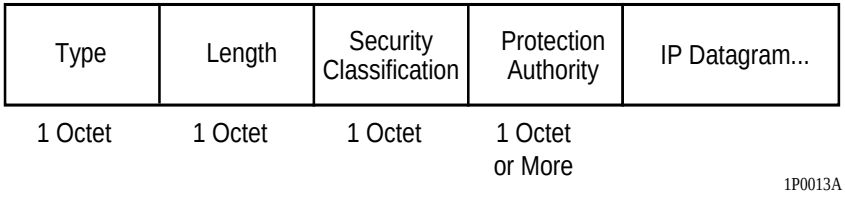

Figure 3-6. RIPSO Security Label

The format of the security label is as follows:

- Octet 1 contains a type value of $82_{(16)}$, identifying the basic security option format.
- Octet 2 specifies the length of the option (three or more octets, depending on the presence or absence of authority flags).

- Octet 3 specifies the security classification levels for the datagrams. Valid security classification levels include

$3D_{(16)}$	Top Secret
$5A_{(16)}$	Secret
$96_{(16)}$	Confidential
$AB_{(16)}$	Unclassified

- Octet 4 and beyond identify the protection authorities under whose rules the datagram is classified at the specified level. (If no authorities have been identified, then this field is not used.)

  The first 7 bits (0 through 6) are flags. Each flag represents a protection authority. The flags defined for Octet 4 are as follows:

Bit 0	GENSER	General Services (as per DoD 5200.28)
Bit 1	SIOP-ESI	DoD (Organization of the Joint Chiefs of Staff)
Bit 2	SCI	Central Intelligence Agency
Bit 3	NSA	National Security Agency
Bit 4	DOE	Department of Energy
Bit 5	Reserved	
Bit 6	Reserved	
Bit 7	Termination indicator	

Note: Bit 7 acts as a “more” bit, indicating that another octet (containing additional authority flags) follows.

Inbound IP Datagrams

When the router receives an IP datagram on a RIPSO interface, it compares the security classification and authority values specified in the security label with those configured on the inbound interface.

If the interface does *not* require a security label for inbound IP datagrams, then the router accepts both unlabeled IP datagrams and datagrams that meet the classification and authority rules described in the next paragraph.

If the interface *does* require a security label, then for the router to accept the datagram, the following RISPO conditions must be met:

- The datagram must be labeled.
- The security classification value in the datagram’s label must be within the security-level range configured for the interface.
- The authority flags in the datagram’s label must include all of the flags required for the interface and cannot contain any flags not allowed for the interface.

The router drops any datagrams that do not meet these requirements and generates an ICMP error message.

On a *non-RIPSO* interface, the router only accepts unlabeled IP datagrams and IP datagrams that are labeled as Unclassified with no authority flags set.

Forwarded IP Datagrams

When the router receives an IP datagram that needs forwarding on a RIPSO interface, the router compares the security classifications and authority values specified in the security label with those configured on the outbound interface. So, before forwarding the datagram, the router

- Checks that all RIPSO conditions are met (see preceding section)
- Applies any outbound-specific configuration parameters

The router drops any datagrams that do not meet these requirements and generates an ICMP error message.

Originated IP Datagrams

When the router originates a datagram and the following conditions are true:

- The datagram needs forwarding through a RIPSO interface
- The RIPSO interface requires outbound labels for originated datagrams

the router labels the datagram with the default security label before transmitting it.

Unlabeled IP Datagrams

If the router receives an unlabeled IP datagram from an interface on which RIPSO is *not* enabled (or on which labels are not required for inbound datagrams), and the IP datagram needs forwarding to an interface on which RIPSO *is* enabled and labels are required for outbound datagrams, then the router labels the datagram, using either an implicit label or default label as follows:

- If the inbound interface has an implicit label configured, then the router uses it to label the datagram.
- If the inbound interface does not have an implicit label configured, then the router labels the datagram with the default label configured for the outbound interface.

If the interface does not have an implicit or default label configured, then the datagram is simply dropped.

RIPSO Example

The router in Figure 3-7 has RIPSO configured on all three IP interfaces. The security ranges specified for each interface vary, as shown. (For simplicity, this example assumes that none of the interfaces requires any authority flags on inbound and outbound traffic, but any flags that are present are acceptable.)

When host 1.1.0.1 broadcasts an all-subnets broadcast IP datagram with the security-level classification set to Secret, the router compares the datagram's classification with the range configured on inbound interface 1.1.0.2. Because Secret is within the range configured on the interface, the router accepts the datagram. In order to forward the datagram, the router does the following:

- Compares the datagram's security level, Secret, to the security-level ranges configured on interface 1.2.0.2 and 1.3.0.2
- Forwards the datagram on interface 1.2.0.2, because Secret is within the security range configured on the interface
- Does *not* forward the datagram on interface 1.3.0.2, because Secret is outside of the security range configured on the interface

Interface	Min. Security Classification	Max. Security Classification
1.1.01	Unclassified	Top secret
1.2.02	Secret	Top secret
1.3.0.2	Top secret	Top secret

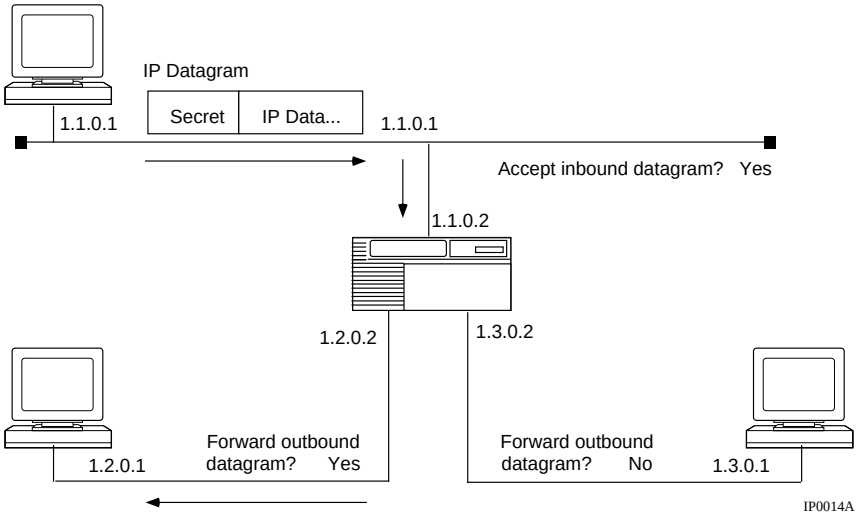

Figure 3-7. RIPSO Network

Defining a Static Route

A *static route* is a manually configured route that specifies the transmission path a datagram must follow, based on the datagram's destination address. A static route specifies a transmission path to another network. You configure a static route if you want to restrict the paths that datagrams follow to paths you specifically configure.

Static routes remain in IP routing tables until you remove them. Note, however, that if the interface that was used to reach the next hop in the static route becomes disabled, the static route disappears from the IP routing table.

You can use Site Manager to

- Specify the state (active or inactive) of the static route record in the IP routing tables.
- Specify the IP address of the network to which you want to configure the static route.
- Specify the subnet mask of the destination network.
- Specify the number of router hops a datagram can traverse before reaching the destination IP address. The IP router uses the cost value when determining the best route for a datagram to follow.
- Specify the IP address of the next-hop router.
- Specify the subnet mask of the next-hop router.
- Specify a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination.
- Specify the local router circuit associated with the static route over an unnumbered interface.

Site Manager: Enable parameter: page A-47
Site Manager: Destination IP Address parameter: page A-48
Site Manager: Address Mask parameter: page A-48
Site Manager: Cost parameter: page A-49
Site Manager: Next Hop Addr parameter: page A-49
Site Manager: Next Hop Mask parameter: page A-49
Site Manager: Preference parameter: page A-50
Site Manager: Unnumbered CCT Name parameter: page A-50

Defining a Black Hole for a Supernet

A router that advertises an aggregate route by using a supernet address to represent multiple explicit routes must be able to discard packets that match the supernet address but that do not match any of the explicit routes.

For example, consider a router that advertises an aggregate route using the supernet address 192.32.0.0/255.255.248. The supernet address represents eight specific networks: 192.32.0.0 through 192.32.7.0. Once the aggregate route has been propagated, the router receives network traffic for each of these specific destinations.

At some point, the router loses connectivity to network 192.32.3.0, one of the networks in the supernet. The router continues to forward traffic that matches destinations 0.0 through 2.0 and 4.0 through 7.0. However, the router can no longer find a complete match in the routing table for the disconnected network, 3.0. The router must drop all traffic destined for 192.32.3.0.

To force the router to drop the packet for an unmatched destination, you configure a special type of static route for a supernet called a *black hole*:

- Enter the supernet address/mask pair as the destination IP address and address mask.
- To create the black hole, enter the black hole encoding (255.255.255.255) as the next-hop address and the next-hop mask.

Configuring Router Discovery

Before a host can send IP datagrams beyond its directly attached subnet, the host must discover the address of at least one operational router on that subnet. Router Discovery is an extension of the Internet Control Message Protocol (ICMP) that enables hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

Routers configured with Router Discovery periodically multicast or broadcast a router advertisement from each of their interfaces, announcing the IP address or addresses of that interface. Hosts discover the addresses of their neighboring routers by listening for these advertisements. Hosts will use the router with the highest preference level as a gateway.

By default, Router Discovery is enabled on each IP interface. You can use Site Manager to select the operating characteristics of Router Discovery on the interface.

Site Manager: Enable parameter: page A-64
Site Manager: Broadcast Type parameter: page A-64
Site Manager: Maximum Interval parameter: page A-65
Site Manager: Lifetime parameter: page A-65
Site Manager: Interface Preference parameter: page A-65

Connecting the Router to a Blacker Front End

The Blacker front end (BFE) is a classified encryption device used by hosts that want to communicate across unsecured wide area networks. BFE devices are typically found in government networks (for example, DSNET), which handle sensitive data requiring a greater degree of security.

Blacker front end support allows the router to connect to BFE devices. The BFE device, in turn, provides the router with encryption services while acting as the data communications equipment (DCE) end of the connection between the router and the X.25 network (Figure 3-8).

Hosts using attached BFE devices can communicate with each other over an unsecured packet-switched network using data paths secured by the encryption services of the BFEs. These hosts are part of a *Red* virtual network. The packet-switched network that carries both the data secured by BFEs and any other unsecured data is known as the *Black* network.

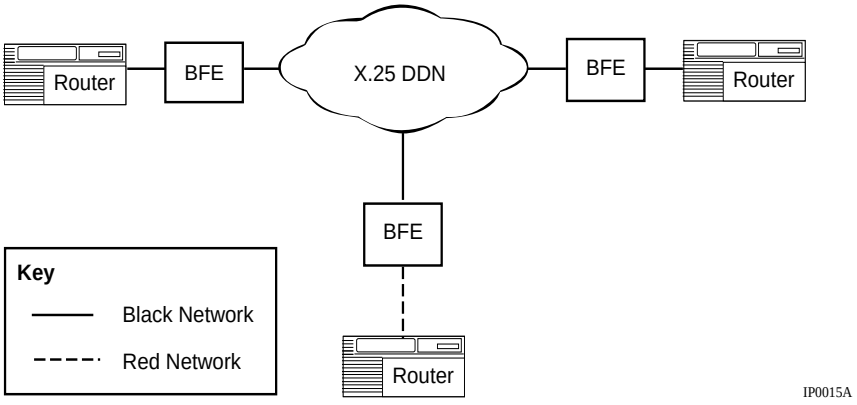

Figure 3-8. Blacker Front End Network Configuration

BFE devices receive authorization and address translation services from an Access Control Center residing on the Black network. The ACC makes access control decisions that determine which hosts are allowed to communicate with each other. A Key Distribution Center (KDC) residing on the Black network provides encryption keys and key management services. A BFE device uses these encryption keys for encrypting traffic between itself and other BFE devices.

The router-to-BFE interface is a modified version of the interface presented in the 1983 DDN X.25 Host Interface Specification. It supports data rates between 1200 b/s and 64 KB/s. In order to support BFE services, the interface must be configured to support IP with the Revised IP Security Option (RIPSO) enabled. All IP datagrams transmitted on the interface must contain a RIPSO security label. The first option in each IP datagram header must be the Basic Security option.

BFE Addressing

You can enable BFE support on individual IP interfaces. When you enable BFE support, the router uses the BFE address-resolution algorithm to map IP addresses to their corresponding X.121 addresses.

BFE IP-to-X.121 address translation differs from standard DDN address translation. Each physical router-to-BFE connection is identified by a BFE X.121 network address and a BFE IP address. The format of a BFE X.121 address is

zzzzzpdddbbb

where:

zzzzz	is zero
p	is the BCD encoding of the port ID
add	is the BCD encoding of the domain ID
bbb	is the BCD encoding of the BFE ID

All BFE hosts are members of Class A IP networks. The format of a BFE IP address is as follows:

nnnnnnnn.Zpppdddd.dddddddb.bbbbbbbb

where:

nnnnnnnn	identifies the network ID in bits
Z	is zero
ppp	is the port ID in bits
dddd.dddddd	is the domain ID in bits
bb.bbbbbbbb	is the BFE ID in bits

BFE supports only physical addressing. It does not support either logical addresses or subaddresses.

Configuring Blacker Front End Support

To configure BFE support on an IP interface, you must

- Configure an X.25 interface that conforms to the BFE requirements described in this section.
- Enable the IP routing protocol on the interface.
- Enable RIPSO support on the interface.

Beginning at the Configuration Manager window, perform the following procedures:

1. **Configure an X.25 interface.**

   When you initially configure packet-level parameters for the X.25 interface, make certain to

   a. Set the Network Address Type parameter to BFE_NETWORK.

   b. Set the DDN IP Address parameter to the IP address that is assigned to your BFE connection.

2. **Edit the packet-layer parameters for the X.25 interface so that they match the settings specified in Table 3-3.**

3. **Add network service record(s) to the X.25 interface.**

4. **Edit the network service record parameters for the X.25 interface so that they match the settings specified in Table 3-4.**

   Remember to set the DDN BFE parameter to Enable.

5. **Enable the IP routing protocol on the X.25 interface.**

   The IP address specified must match the one specified in the packet-layer parameter setting.

6. **Edit the IP interface record.**

   The address resolution must be set to X.25 BFE DDN. Also configure IP security options (RIPSO) on the interface. IP security must be enabled and labels are required on all outbound data.

For instructions on performing Steps 1 through 4, see *Configuring X.25 Services*. For instructions on performing Steps 5 and 6, see the section "Configuring the Revised IP Security Option on an Interface" on page 3-37.

Note: Generally, the synchronous line parameter settings are the same for both a DDN X.25 link and a BFE X.25 link. However, if your operating environment has specific needs, you may want to edit synchronous line parameters. See the appropriate protocol manual for instructions.

Table 3-3. BFE X.25 Packet-Level Parameter Settings

X.25 Parameter	BFE Required Setting
Enable	Enable
Network Address Type	BFE_NETWORK
PDN X.121 Address	Parameter is ignored.
DDN IP Address	Specify the IP address assigned to your BFE connection.
Sequence Size	MOD8
Restart Procedure Type	DTE_RESTART
Default Tx/Rx Window Size	BFE range is 2 to 7. This setting should match the default value configured in the BFE. This value should be coordinated with the X.25 service record value.
Default Tx/Rx Packet Length	BFE options include 128, 256, 512, and 1024. This setting should match the default value configured in the BFE.This value should be coordinated with the X.25 service record value.
Number of incoming SVC channels	Zero (0). BFE does not support the one-way logical channel incoming facility.
Incoming SVC LCN Start	Parameter is ignored.
Number of outgoing SVC channels	Any valid nonzero setting.
Bidirectional SVC LCN	Any valid nonzero setting.
Number of outgoing SVC channels	Zero (0). BFE does not support the one-way logical channel outgoing facility.
Outgoing SVC LCN Start	Parameter is ignored.
Number of PVC channels	Zero (0). BFE does not support PVCs.
PVC LCN Start	Parameter is ignored.
T1 Timer, T2 Timer, T3 Timer, T4 Timer	BFE has no special requirements for any of these four parameters.
Flow Control Negotiation	Set to On if you do not want to use the default values configured in the BFE for this link.

(continued)

Table 3-3. BFE X.25 Packet-Level Parameter Settings *(continued)*

X.25 Parameter	BFE Required Setting
Max Window Size	BFE range is 2 to 7. If you specify any other setting than the default value configured in the BFE, set Flow Control Negotiation to On.This value should be coordinated with the X.25 service record value.
Max Packet Length	BFE options include 128, 256, 512, and 1024. If you specify any other value than the default value configured in the BFE, then set Flow Control Negotiation to On. (If the IP interface is configured to support multiple IP security levels, then set to 1024.) This value should be coordinated with the X.25 service record value.
Trans/Recv Throughput Class	Parameter is ignored.
Max Throughput Class	Parameter is ignored.
Throughput Class Negotiation	Off
Network User Identification	Off
Incoming Calls Accept	On
Outgoing Calls Accept	On
Fast Select Accept	Off
Reverse Charge Accept	Off
Fast Select	Off
Reverse Charging	Off
CUG Selection	Null
CUG Outgoing Access	Null
CUG Bilateral Selection	Null
RPOA Selection	Off
Charging Information	Off
Transit Delay	Off
Full Addressing	On
Acceptance Format	Defext
Release Format	Defext
CCITT (now ITU-T) Conformance	DXE1980
Network Standard	DOD

Table 3-4. BFE X.25 Network Service Record Parameter Settings

X.25 Parameter	BFE Required Setting
Enable	Enable
Type	DDN
Connection ID	Parameter is ignored.
Remote IP Address	Specify the IP address of the remote system.
Remote X.121 Address	Parameter is ignored.
Broadcast	Parameter is ignored.
Max Connections	Any valid setting.
Precedence	Any valid setting. The BFE will accept, but not act on, the DDN Precedence facility.
Max Idle	Any valid setting.
Call Retry	Any valid setting.
Flow Facility	Set to On if you want to use a value other than the default window size and packet size configured in the BFE.
Window Size	BFE range is 2 to 7. If you want to use a value other than the default window size configured in the BFE, set Flow Facility to On. You must coordinate this value with the packet-level value.
Packet Size	BFE options include 128, 256, 512, and 1024. If you want to use a value other than the default packet size configured in the BFE, set Flow Facility to On. (If IP interface is configured to support multiple IP security levels, then set to 1024.) You must coordinate this value with the packet-level value.
Fast Select Request	Off
Fast Select Accept	Off
Reverse Charge Request	Off
Reverse Charge Accept	Off
User Facility	Null
DDN BFE	Enable
CUG Facility Format	None

(continued)

Table 3-4. BFE X.25 Network Service Record Parameter Settings

X.25 Parameter	BFE Required Setting
CUG Facility Type	This parameter is ignored.
CUG Number	This parameter is ignored.

Chapter 4
Customizing RIP Services

This chapter consists of the following sections:

- Enabling and Disabling RIP on an Interface -- page 4-2
- Selecting the RIP Version -- page 4-2
- Sending RIP Updates on an Interface -- page 4-3
- Sending Triggered Updates -- page 4-4
- Specifying a Time-to-Live Value -- page 4-5
- Receiving RIP Updates on an Interface -- page 4-5
- Authenticating the Password on a Version 2 Update -- page 4-6
- Supplying a Default Route on an Interface -- page 4-7
- Listening for a Default Route -- page 4-7
- Configuring RIP for Equal-Cost Multipath Support -- page 4-8
- Configuring a RIP Interface for Dial-Optimized Routing -- page 4-9
- Setting RIP Timers on an Interface -- page 4-9

  Specifying an Update Interval -- page 4-9
  Specifying a Timeout Period -- page 4-10
  Specifying a Holddown Period -- page 4-10
  Specifying a Stabilization Time -- page 4-11

Enabling and Disabling RIP on an Interface

When you add RIP to an IP interface, RIP is automatically enabled on that interface.

You can use Site Manager to disable and re-enable RIP on an interface.

Site Manager: Enable parameter: page A-94

Selecting the RIP Version

You can specify whether RIP sends Version 1 updates, Version 2 updates with no aggregation of subnets, or Version 2 updates with subnet aggregation.

- In RIP Version 1 mode, RIP generates Version 1 updates only, using the broadcast address as specified in RFC 1058. RIP aggregates subnet information.
- In RIP Version 2 mode, RIP generates Version 2 updates, using the multicast address 224.0.0.9 as specified in the RIP Version 2 RFC, 1388. RIP does not aggregate subnet information.
- In RIP Version 2 mode with aggregation, RIP generates RIP Version 2 updates, using the multicast address, and performs aggregation of subnets into a natural network advertisement on interfaces belonging to another network.

In both Version 2 modes, RIP checks for a password on all received updates (see “Authenticating the Password on a Version 2 Update” on page 4-6).

By default RIP sends Version 1 updates. You can use Site Manager to select Version 2 with or without aggregation of subnet information.

Site Manager: RIP Mode parameter: page A-98

Sending RIP Updates on an Interface

By default, RIP supplies RIP updates to neighboring networks on each interface. You can use Site Manager to disable and re-enable this feature.

Site Manager: RIP Supply parameter: page A-94

RIP can issue routing updates in the following modes:

- Poisoned reverse updates
- Split horizon updates
- Actual cost updates

Poisoned reverse updates and split horizon updates are schemes for controlling the way a router advertises a route to the neighbor from which it learned the route.

In *poisoned reverse* updating, a router that sends updates to a neighbor includes routes learned from that neighbor but sets the route metric to infinity.

In *split horizon* updating, a router that sends updates to a neighbor omits routes that it learned from that neighbor.

On certain interfaces -- for example, on a Frame Relay interface that has virtual connections (VCs) to different routers that are part of the same logical IP subnet -- you may need to advertise all learned routes with the learned cost.

By default, RIP sends poisoned reverse updates. Using Site Manager, you can configure RIP to send split horizon and actual cost updates.

Site Manager: Poisoned Reverse parameter: page A-96

Sending Triggered Updates

A RIP router generates full routing updates at regular intervals. You can also configure RIP to generate an update on a specified interface each time it recalculates a route's metric. Such an update is called a *triggered update*. A triggered update contains only the routes that have changed. (RIP also sends full updates at regular intervals on interfaces configured for triggered updating.)

By default, triggered updates are disabled on all interfaces configured for RIP.

You can use Site Manager to enable and disable triggered updates on this interface.

Site Manager: Triggered Updates parameter: page A-99

For compatibility with routers running Version 8.10 or earlier, disable this feature. Pre-Version 9.00 implementations of RIP do not support triggered updates.

Specifying a Time-to-Live Value

By default, RIP inserts a time-to-live (TTL) value of 1 hop into each outbound routing update. Setting a TTL of 1 prevents RIP updates from inadvertently getting off the local network. Increasing the TTL introduces the risk of the update getting off the local network and being forwarded around the network.

Certain RIP implementations ignore packets with a TTL value of 1 hop. Use this parameter to provide interoperability with such implementations.

You can use Site Manager to specify a TTL value greater than 1.

Site Manager: Time to Live parameter: page A-97

Receiving RIP Updates on an Interface

By default, RIP listens for routing updates on every interface on which it is enabled and stores the route in the routing table.

You can use Site Manager to disable this feature.

Site Manager: RIP Listen parameter: page A-95

Authenticating the Password on a Version 2 Update

By default, RIP running in RIP Version 2 mode does not authenticate the password on incoming updates. RIP checks for the presence of a password:

- If no password is present, RIP accepts the update.
- If a password is present, RIP drops the update.

With authentication enabled, RIP drops all received Version 1 updates and processes only Version 2 updates in the following manner:

- If no password is present in the Version 2 update, RIP drops the update.
- If a password is present in the Version 2 update and that password is valid, RIP accepts the update.
- If the password is invalid, RIP drops the update.

Using Site Manager, you can configure a RIP interface for authentication and enable password checking.

If you configure authentication on an IP interface, you can use Site Manager to assign a 1-to-16 character password to the interface.

Site Manager: Authentication Type parameter: page A-99
Site Manager: Authentication Password parameter: page A-99

Supplying a Default Route on an Interface

When the routing table does not contain the route to a particular destination address, the router looks for a default route to the destination. Like any other route in the routing table, the default route can be acquired dynamically (by means of a routing protocol) or entered statically (by you).

By default, RIP does not supply a default route. You can use Site Manager to enable this feature on a RIP interface.

You can configure RIP to advertise an existing default route (one that is present in the routing table) in RIP updates sent to neighboring networks. You can also configure RIP to generate a default route if the routing table does not contain a default route.

Site Manager: Default Route Supply parameter: page A-95

This parameter is independent of the RIP Supply parameter. A configured policy can override the parameter.

Listening for a Default Route

By default, RIP ignores inbound advertisements of a default route on the interfaces where it is configured. Using Site Manager, you can configure RIP to listen for a default route.

Site Manager: Default Route Listen parameter: page A-96

Configuring RIP for Equal-Cost Multipath Support

By default, the IP routing table contains a single “best” RIP route to a given destination. If RIP submits another route to the same destination, IP compares the new route with the current route. If the new route is better, IP replaces the current route with the existing new route. If not, IP discards the new route.

If you have enabled equal-cost multipath support on the router, IP can store multiple equal-cost best routes in the routing table. When RIP submits a route to a destination, one of the following events occurs:

- IP determines that the current route to that destination is better than the new route. IP discards the new route.
- IP determines that the new route is better than the current route. IP discards the current route and replaces it with the new route. In the event that the routing table contains multiple equal-cost best routes, IP discards all of these routes.
- IP determines the new route and the current route have the same cost. IP adds the new route to the routing table -- up to a maximum number that you specify. If the routing table already contains the maximum number of equal-cost routes learned from RIP, IP discards the routes.

You can use Site Manager to specify the number of equal-cost routes (up to five) that IP can store in the routing table.

Site Manager: RIP Maximum Equal Cost Paths parameter: page A-45

For information about configuring equal-cost multipath support on the router, see “Enabling Equal-Cost Multipath Support” on page 3-11.

Configuring a RIP Interface for Dial-Optimized Routing

Dial-optimized routing is a method of reducing costs on dialed lines. Under dial-optimized routing, RIP exchanges routing information on the interface only when the router or a peer has activated the connection for a data transmission. RIP does not initiate a connection on a dialed line solely for the purpose of issuing a routing update.

For information on enabling dial-optimized routing, see *Configuring Dial Services.*

Once you have enabled dial-optimized routing, you can use Site Manager to set various RIP timers to control the way RIP generates updates on interfaces to dialed lines. For information on setting timers for dial-optimized routing, see "Setting RIP Timers on an Interface" on page 4-9.

Setting RIP Timers on an Interface

Configurable timers determine the way RIP manages route information on an interface. Setting these timers allows you to specify

- The frequency at which RIP broadcasts full RIP updates on the interface
- The timeout period that RIP will wait before considering a network unreachable
- The holddown period that unreachable routes will be retained in the routing table and advertised

Specifying an Update Interval

By default, RIP generates a full update every 30 seconds on each interface configured with RIP.

If you have enabled dial optimized routing on this interface, the default is 1 hour.

You can use Site Manager to specify an update interval for the interface.

Site Manager: Broadcast Timer parameter: page A-97

Specifying a Timeout Period

By default, RIP waits 90 seconds for an update from a network before it considers that network to be unreachable.

If you have enabled dial optimized routing on this interface, the default is 3 hours.

You can use Site Manager to specify a timeout period between 15 seconds and 259,200 seconds (72 hours). For dial optimized routing, the maximum value is 3,628,800 seconds (6 weeks).

Site Manager: Timeout Timer parameter: page A-97

Specifying a Holddown Period

Once RIP has determined that a network is unreachable, RIP continues to advertise a route to that network for a default holddown period of 90 seconds.

If you have enabled dial optimized routing on this interface, the default is 3 hours.

You can use Site Manager to specify a holddown period between 15 seconds and 259,200 seconds (72 hours). For dial optimized routing, the maximum value is 3,628,800 seconds (6 weeks).

Site Manager: Holddown Timer parameter: page A-98

Specifying a Stabilization Time

The stabilization time is the period that RIP allows itself to learn all routes from its neighbors before sending full updates. By default, RIP uses a stabilization time of 120 seconds.

You can use Site Manager to specify a stabilization time for this interface.

Site Manager: Stabilization Timer parameter: page A-100

Chapter 5
Customizing OSPF Services

OSPF (Open Shortest Path First) is an internal gateway protocol for use in large networks. This chapter consists of the following sections that describe OSPF and show you how to edit OSPF parameters.

- OSPF Concepts and Terminology -- page 5-3
  - OSPF Addresses and Variable-Length Masks -- page 5-3
  - OSPF Neighbors -- page 5-3
  - Neighbor Adjacencies -- page 5-4
  - Designated Routers -- page 5-4
  - OSPF Areas -- page 5-5
  - Boundary Routers and AS External Routes -- page 5-6
- Configuring OSPF Global Features -- page 5-9
  - Enabling and Disabling OSPF on the Router -- page 5-9
  - Supplying an OSPF ID -- page 5-9
  - Configuring the Soloist and Backup Soloist on a Slot -- page 5-10
  - Using the Holddown Timer to Relieve the CPU -- page 5-11
  - Configuring Message Logging -- page 5-11
  - Configuring OSPF for Equal-Cost Multipath Routing -- page 5-13

- Configuring OSPF on an IP Interface -- page 5-14

  Enabling and Disabling OSPF -- page 5-14
  Configuring an Area ID -- page 5-14
  Specifying the Network Type -- page 5-15
  Specifying Router Priority for a Multiaccess Network -- page 5-17
  Estimating the Transit Delay -- page 5-18
  Configuring Interface Timers -- page 5-19
  Specifying Metric Cost -- page 5-23
  Specifying the MTU Size -- page 5-24
  Configuring OSPF on a Circuitless Interface -- page 5-25
  Configuring a Neighbor on an NBMA Interface -- page 5-25

- Configuring a Boundary Router -- page 5-25

  Enabling the Boundary Function -- page 5-25
  Configuring AS External Advertisements -- page 5-26
  Using the External Route Tag in an ASE -- page 5-28

- Configuring Authentication on an Interface -- page 5-30
- Configuring an Area Border Router -- page 5-30

  Configuring a Virtual Backbone Link through a Transit Area -- page 5-32
  Defining a Range for a Summary Route -- page 5-34
  Configuring Border Services for a Stub -- page 5-34

OSPF Concepts and Terminology

OSPF is a *link-state protocol*. A router running a link-state protocol periodically tests the status of the physical connection to each of its neighbor routers and sends this information to its other neighbors. A link-state protocol does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information throughout the autonomous system (or area, if the AS is divided into areas). This process is referred to as the synchronization of the routers' topological databases.

With the link information, each router builds a *shortest-path tree* with itself as the root of the tree. It then can identify the shortest path from itself to each destination and build its routing table.

OSPF Addresses and Variable-Length Masks

A destination in as OSPF route advertisement is expressed as an IP address and a variable-length mask. Taken together, the address and the mask indicate the range of destinations to which the advertisement applies.

The ability to specify a range of networks allows OSPF to send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 0xffff0000 describes a single route to destinations 128.185.0.0 to 128.185.255.255.

OSPF Neighbors

OSPF *neighbors* are any two routers that have an interface to the same network. In each OSPF network, routers use the Hello protocol to discover their neighbors and maintain neighbor relationships. On a broadcast or point-to-point network, the Hello protocol dynamically discovers neighbors; however, on a nonbroadcast multiaccess network, you must manually configure neighbors.

The Hello protocol is responsible for ensuring that communication between neighbors is bidirectional. Periodically, OSPF routers send out Hello packets over all interfaces. Included in these Hello packets are

- The router's priority
- The router's Hello Timer and Dead Timer value
- A list of routers that have sent this router Hello packets on this interface
- The router's choice for designated router and backup designated router

Bidirectional communication is determined when one router sees itself listed in the neighbor's Hello packet.

Neighbor Adjacencies

Neighbors may form an *adjacency* for the purpose of exchanging routing information. When two routers form an adjacency, they go through a process called *database exchange* to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. From this point on, only routing change information is passed between the adjacencies, thus conserving bandwidth.

All routers connected by a point-to-point network or a virtual link will always form an adjacency. Also, every router on a multiaccess network forms an adjacency relationship with the designated router and the backup designated router.

Designated Routers

To further reduce the amount of routing traffic, the Hello protocol elects a *designated router* and a *backup designated router* on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form an adjacency with the designated router and the backup designated router only and send link state information to them. The designated router then redistributes the information from each router to every other router.

The Hello protocol always elects a backup designated router along with the designated router. Its responsibility is to take over all of the designated router's functions should the designated router fail.

OSPF Areas

OSPF routers reduce and restrict the amount of internal and external routing information that is flooded through the AS by dividing the AS into *areas*. Figure 5-1 shows an OSPF autonomous system divided into three areas and a required central area called a *backbone* that is used to distribute routing information among areas.

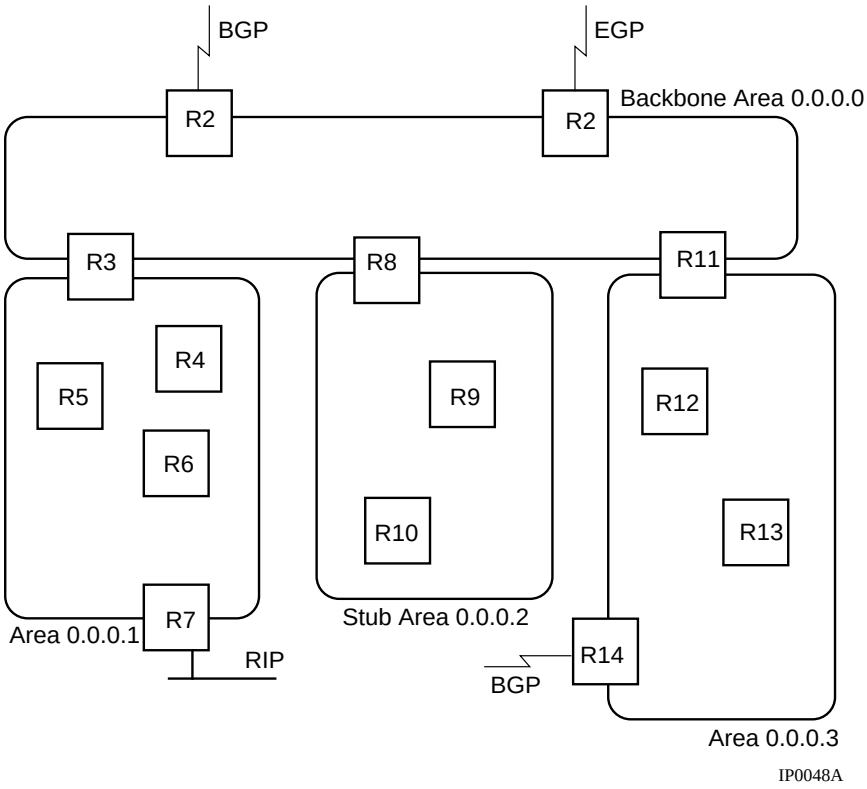

Figure 5-1. OSPF Areas

Area Internal Routers

A router with interfaces to networks in one area only is considered to be an area *internal router.* Internal routers flood each area with complete routing information about changes that occur within the area. In Area 0.0.0.1, for example, R4, R5, and R6 are internal routers.

Area Border Router

A router with an interface to the backbone network and interfaces to one or more additional areas is considered to be a *border router.* Each border router connects one or more areas to the backbone. In Figure 5-1, R3, R8, and R11 are border routers. Using the backbone, border routers ensure that ASEs and summaries of routing information for all areas are distributed throughout the AS.

For further information and instructions, see “Configuring an Area Border Router” on page 5-30.

Area Identifiers

Each area has a unique ID number. (ID 0.0.0.0 is always reserved for the OSPF backbone.) The AS in Figure 5-1 has three areas (0.0.0.1, 0.0.0.2, and 0.0.0.3) and a backbone (0.0.0.0).

For information on associating an OSPF interface with an area ID, see “Configuring an Area ID” on page 5-14.

Boundary Routers and AS External Routes

OSPF considers the following routes to be *AS external (ASE)* routes:

- A route to a destination outside the AS
- A static route
- A default route
- A route derived by RIP
- A directly connected network not running OSPF

A router configured with BGP, RIP, or another protocol to receive information about external routes and OSFP to inject this information into an OSPF AS is considered to be an OSPF *boundary router.*

In Figure 5-1, for example, routers R1 and R2 are boundary routers that use BGP and EGP to connect the backbone to external ASs. R7 in Area 0.0.0.1 is also a boundary router, connecting the area to an external RIP network. R14 in Area 0.0.0.3 connects the area to an external AS via BGP.

For further information and instructions, see “Configuring a Boundary Router” on page 5-25.

OSPF Implementation Notes

This section provides some suggestions to help you configure your OSPF network. The Bay Networks OSPF implementation does not restrict you to these suggestions, but we are providing them as guidelines.

- Keep the same password throughout an area, or even throughout the entire OSPF AS, if possible.
- Use the default timers, unless you are running 9.6-KB sync lines. In this case, double the default timers on both ends of the link.
- Use address ranges if your network is a subnetted network.
- Keep all subnets within one area. If you cross areas, you cannot configure summaries.
- Make sure the AS Boundary Router function is enabled if the router has any non-OSPF interfaces and if you want that information propagated.
- You must configure virtual links for each area border router that does not reside within or directly interface to the backbone. Every area border router must have a configured path to the backbone.
- Rather than just a hop count, OSPF considers the cost of a path when choosing the best path. Each interface, however, is assigned the default cost 1 for the path to which it interfaces. If you have a preferred path, you must edit the Metric Cost parameter for your interfaces. You will need to assign a higher metric cost for those paths which are *not* preferred paths.
- If you have any devices in your network running OSPF, and are now adding a Bay Networks router, you must make sure that the router’s timer values coincide with the timers in your other devices. Determine the timer values of the other devices, and change the router’s timer values to match them.

- If you change the topology (for example, if you add an area, combine two areas, move routers, and so on), you must reconfigure the appropriate OSPF elements (OSPF area ranges/interfaces/neighbors/virtual links, and so on).

Configuring OSPF Global Features

OSPF global features affect the way OSPF runs on the router. They apply to all OSPF interfaces.

Enabling and Disabling OSPF on the Router

When you start OSPF on the router, OSPF is automatically enabled.

You can use Site Manager to disable and re-enable OSPF on the router.

Site Manager: Enable parameter: page A-73

Supplying an OSPF ID

Each router configured with OSPF has an OSPF ID. This IP address uniquely identifies this router in the OSPF domain. By default, OSPF uses the IP address of the first OSPF circuit configured on this router.

By convention, and to ensure uniqueness, the router ID should be one of the router's IP interface addresses.

You can use Site Manager to supply an OSPF ID (overriding the default).

Site Manager: Router ID parameter: page A-73

The router ID determines the designated router on a broadcast link if the priority values of the routers being considered are equal. The higher the router ID, the greater its priority.

If both OSPF and BGP are running on the router, the OSPF router ID must be identical to the BGP identifier. In addition, the OSPF router ID must match one of the IP addresses configured on the router.

Configuring the Soloist and Backup Soloist on a Slot

The OSPF protocol is implemented as a *soloist* -- that is, as a single process running on a single slot of a router. When you add an OSPF interface to a circuit, the router enables OSPF on a slot. If the slot on which the OSPF soloist is running goes down, the router will attempt to run OSPF on another slot.

Each time the OSPF soloist is restarted, all of the routing information is lost and must be relearned from the network. The *OSPF backup soloist* provides a method of preserving information learned from the network in the event of an OSPF crash or slot removal, avoiding the time-consuming and resource-intensive process of relearning routing information. In the event of a crash or slot removal, transition between the OSPF primary and backup soloist occurs without relearning routing information from the network.

By default, the router uses any available slot for the OSPF soloist. You can use Site Manager to specify the slot or slots or which the OSPF soloist runs.

Site Manager: OSPF Slot parameter: page A-75

By default, OSPF does not maintain a copy of the link state database (LSDB) for the backup soloist. You can use Site Manager to maintain an LSDB for the backup.

Site Manager: Backup Enable parameter: page A-75

Using the Holddown Timer to Relieve the CPU

The holddown timer controls how often OSPF performs the calculation to compute a route. Its purpose is to free up the CPU. Note that a value of 0 means there is no holddown time.

By default, the holddown timer is set at 1 second.

You can use Site Manager to specify a holddown time.

Site Manager: Hold Down Timer parameter: page A-74

Configuring Message Logging

OSPF maintains a primary and backup log for OSPF messages.

By default, OSPF logs the following messages in the primary log:

TRACE
INFO
DEBUG
INTF STATE
NBR STATE
BAD LS

By default, OSPF logs no messages in the backup log.

You can use Site Manager to specify the messages that OSPF writes to the primary and backup log.

Site Manager: Primary Log Mask parameter: page A-76
Site Manager: Backup Log Mask parameter: page A-76

OSPF can log any or all of the following messages:

Message	Example
TRACE	"Designated Router changed on network :*x.x.x.x x.x.x.x* -> *x.x.x.x*"
INFO	"OSPF enabled"
OSPF DEBUG	"OSPF couldn't get a buffer, dying"
INTF STATE	"Interface *x.x.x.x* up on circuit *x*"
NBR STATE	"T2: Neighbor *x.x.x.x* Event: *x* State change: *x*->*x*"
LSA SELF-ORIGIN	"T4: Originating new LSA - type *x* LSID *x.x.x.x* router *x.x.x.x*"
LSA RECEIPT	"T5: Received new LSA - type *x* LSID *x.x.x.x* router *x.x.x.x* neighbor *x.x.x.x*"
ROUTE CHANGE	"T6: Routing Table changed - type *x* destination *x.x.x.x* old *x.x.x.x* new *x.x.x.x*"
BAD LS	"R4: Ack received for non-existent LSA: type *x* LSID *x.x.x.x* neighbor *x.x.x.x*"
LESS RECENT LSA	"C3: Packet Rejected: LS UPDATE: LESS RECENT RX (*x*) src *x.x.x.x* type *x* ls_id: *x.x.x.x* adv_rtr: *x.x.x.x* ls_seq: *x* ls_age: *x* db_seq: *x* db_age: *x* elapse: *x* freeme:*x* ackcnt:x nbr_retrans:x nbrEcnt:x Fcnt:x"
MORE RECENT LSA	"R3: Received more recent self-originated LSA: type *x* LSID *x.x.x.x* router *x.x.x.x* neighbor *x.x.x.x*"
MAX AGE LSA	"N3: LSA of MaxAge flushed: type x LSID *x.x.x.x* router *x.x.x.x*"

Configuring OSPF for Equal-Cost Multipath Routing

By default, the IP routing table contains a single “best” OSPF route to a given destination. If OSPF submits another route to the same destination, IP compares the new route with the current route. If the new route is better, IP replaces the current route with the existing new route. If not, IP discards the new route.

If you have enabled equal-cost multipath support on the router, IP can store multiple equal-cost best routes in the routing table. When OSPF submits a route to a destination, one of the following events occurs:

- IP determines that the current route to that destination is better than the new route. IP discards the new route.
- IP determines that the new route is better than the current route. IP discards the current route and replaces it with the new route. In the event that the routing table contains multiple equal-cost best routes, IP discards all of these routes.
- IP determines that the new route and the current route have the same cost. IP adds the new route to the routing table -- up to a maximum number that you specify. If the routing table already contains the maximum number of equal-cost routes learned from OSPF, IP discards the routes.

You can use Site Manager to specify the number of equal-cost routes (up to five) that IP can store in the routing table.

Site Manager: OSPF Maxium Path parameter: page A-77

For information about configuring equal-cost multipath support on the router, see “Enabling Equal-Cost Multipath Support” on page 3-11.

Configuring OSPF on an IP Interface

The following sections describe features that you can configure on an OSPF interface.

Enabling and Disabling OSPF

When you add OSPF to an IP interface, OSPF is automatically enabled on that interface. The interface will be advertised as an internal route. In addition, the interface can be used to form a neighbor relationship.

You can use Site Manager to disable and re-enable OSPF on an interface.

Site Manager: Enable parameter: page A-78

Each OSPF interface on the router provides access to one area in the OSPF domain. By default, OSPF assumes that each interface provides access to the backbone area.

Configuring an Area ID

In an AS that contains multiple areas, each OSPF interface is configured with the ID of the area to which it is connected.

In Figure 5-1, for example, R5 is an internal router in Area 0.0.0.1. Each of its interfaces is attached to a network in that area. Therefore, each interface is configured with area ID 0.0.0.1.

By default, OSPF assigns area ID 0.0.0.0 (the backbone ID) to each interface.

You can use Site Manager to assign an area ID to an OSPF interface.

Site Manager: Area ID parameter: page A-78

Specifying the Network Type

By default, OSPF assumes that the interface is attached to a broadcast network.You can use Site Manager to specify one of the following network types:

Broadcast	Default. Select Broadcast if this network is a broadcast LAN, such as Ethernet.
NBMA	Select nonbroadcast multiaccess if the network is a nonbroadcast network, such as X.25.
Point-to-point	Select Point-to-point for a synchronous, point-to-point interface.
Point-to-multipoint (STD)	Select Point-to-multipoint (STD) if the network is a point-to-multipoint network.
Point-to-multipoint	Select Point-to-multipoint if you want to use the Bay Networks proprietary solution for Frame Relay point-to-multipoint networks.

Note: If you set this parameter to NBMA, you need to configure neighbors manually.

Site Manager: Type parameter: page A-79

A *broadcast network* supports multiple routers and can address a single physical message to all attached routers. Examples of such a network are Ethernet, FDDI, and Token Ring.

A *nonbroadcast multiaccess* (NBMA) network supports multiple routers and cannot address a single physical message to all routers. Examples of such a network are Frame Relay and X.25.

A *point-to-point network* joins a single pair of OSPF routers. An example of such a network would be a network of synchronous lines.

A *point-to-multipoint network* supports multiple routers in a partial mesh configuration. (Bay Networks supports the standard OSPF point-to-multipoint interface and also provides a proprietary point-to-multipoint solution for routers running OSPF in star Frame Relay topologies.)

Figure 5-2, for example, shows a point-to-multipoint topology in which four AN routers are connected by Frame Relay links to a BCN router. The AN routers are the spokes of the topology, and the BCN router is the hub. All of the routers are running OSPF. The BCN router is connected to the Frame Relay network over a permanent virtual circuit (PVC) in group mode. The AN routers are connected over PVCs in direct or group mode. For details on Frame Relay, see *Configuring Frame Relay Services.*

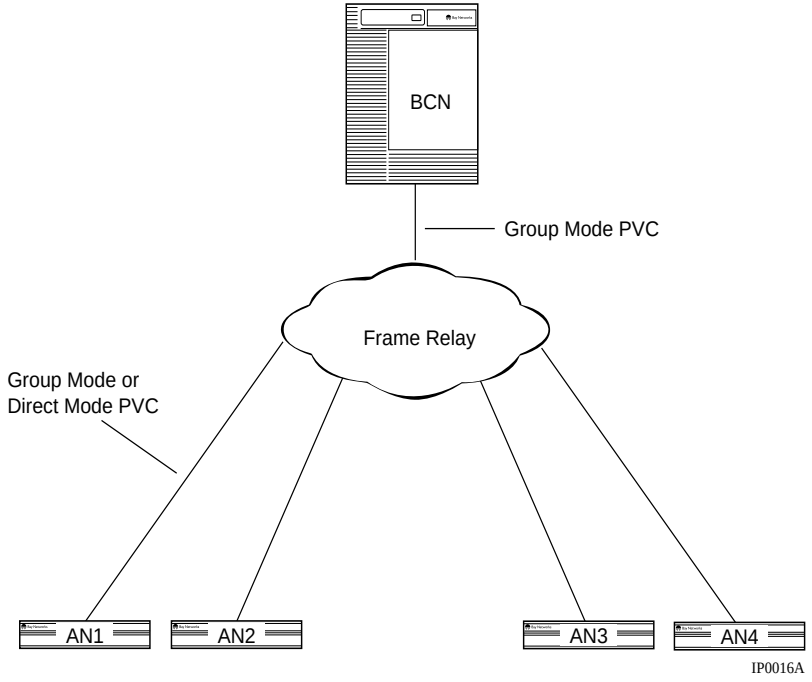

Figure 5-2. Point-to-Multipoint Topology

OSPF point-to-multipoint interfaces provide an efficient means to connect routers in a star topology. The routers are configured as follows:

- The hub of the star topology -- the BCN router in Figure 5-2 -- is configured with a point-to-multipoint interface to the PVC and is set to be the OSPF designated router in the network. The Router Priority parameter is set to a value greater than 0.
- Each spoke of the star -- the AN routers in Figure 5-2 -- is configured with a point-to-multipoint interface to the PVC and is made ineligible to become the designated router. The Router Priority parameter on each AN is set to 0.

When the spokes of the topology (the AN routers) are computing routes through the other spokes, the next hop is forced to be the hub (the BCN router). The hub can then forward the packet to the correct spoke.

Running OSPF with point-to-multipoint network interfaces addresses two problems: how to minimize the number of subnets and how to minimize the number of interfaces required to support communications within the star topology. With point-to-multipoint interfaces, each star topology requires only one subnet, rather than one subnet for each PVC. Also, the hub needs to support only one interface for each star rather than one interface for each PVC. This reduces the demand for resources on the router.

Specifying Router Priority for a Multiaccess Network

The router priority value is used in multiaccess networks (Broadcast, NBMA, or point-to-multipoint) to elect the designated router.

A router with a priority of 0 is not eligible to become the designated router on this particular network.

By default, each OSPF interface has a router priority of 1.

You can use Site Manager to

- Specify a priority value for the interface
- Make the router ineligible to be a designated router on this interface

Site Manager: Rtr Priority parameter: page A-79

In the case of equal router priority values, the router ID will determine which router will become the designated router. However, if there already is a designated router on the network when you start this router, it will remain the designated router no matter what your priority or router ID.

Estimating the Transit Delay

By default, OSPF estimates that it takes one second to transmit a link-state update packet over this interface.

You can use Site Manager to supply a different transit delay estimate.

Site Manager: Transit Delay parameter: page A-80

Configuring Interface Timers

The following sections describe the timers that you can set on each OSPF interface.

Setting the Retransmit Interval

The retransmit interval is the number of seconds between link-state advertisement retransmissions for adjacencies belonging to this interface.

Each type of network has an optimum retransmit interval. If the interface is connected to a broadcast network, we suggest you use the default setting -- 5 seconds. We suggest the following settings for network types supported by OSPF:

Network Type	Suggested Retransmit Interval
Broadcast	5 seconds (default)
Point-to-point	10 seconds
NBMA	10 seconds
Point-to-multipoint	10 seconds

You can use Site Manager to specify a retransmit interval.

Site Manager: Retransmit Interval parameter: page A-80

Setting the Hello Interval

The hello interval specifies how often the router sends hello messages on the interface. By default, OSPF transmits a hello message every 10 seconds.

Each type of network has an optimum hello interval. If the interface is connected to a broadcast network, we suggest you use the default setting -- 10 seconds. We suggest the following settings for network types supported by OSPF:

Network Type	Suggested Hello Interval
Broadcast	10 seconds (default)
Point-to-point	15 seconds
NBMA	20 seconds
Point-to-multipoint	15 seconds

You can use Site Manager to specify a hello interval.

Note: This value must be the same for all routers attached to the same network.

Site Manager: Hello Interval parameter: page A-81

Setting the Dead Interval

The dead interval is the number of seconds that a router waits to receive a Hello packet from a neighbor before considering the neighbor to be down. The Dead Interval value should be some multiple of the Hello Interval value.

Each type of network has an optimum dead interval. If the interface is connected to a broadcast network, we suggest you use the default setting -- 40 seconds. We suggest the following settings for other network types supported by OSPF:

Network Type	Suggested Dead Interval
Broadcast	40 seconds (default)
Point-to-point	60 seconds
NBMA	80 seconds
Point-to-multipoint	60 seconds

You can use Site Manager to specify a dead interval.

Note: This value must be the same for all routers attached to the same network.

Site Manager: Dead Interval parameter: page A-81

Poll Interval for NBMA Neighbors

The poll interval is the largest number of seconds allowed between Hello packets sent to an inactive nonbroadcast multiaccess neighbor.

By default, each OSPF interface has a poll interval of 120 seconds.

You can use Site Manager to specify a poll interval.

Site Manager: Poll Interval parameter: page A-82

Specifying Metric Cost

For OSPF, the best path is the one that offers the least-cost metric.

By default, each OSPF interface has a cost of 1. You can use Site Manager to specify a metric cost for the interface.

Site Manager: Metric Cost parameter: page A-82

You must configure cost metrics if you want to specify a preferred path. Allow the preferred path to retain the cost metric value of 1, then assign higher-cost metric values to the less-preferred paths.

Figure 5-3 shows the benefit of using configurable cost metrics. Assigning the 56-Kb line a cost metric value of 10 forces OSPF to choose the faster T1 line path as the best path, despite the extra hop, when transmitting a packet from Host A to Host B.

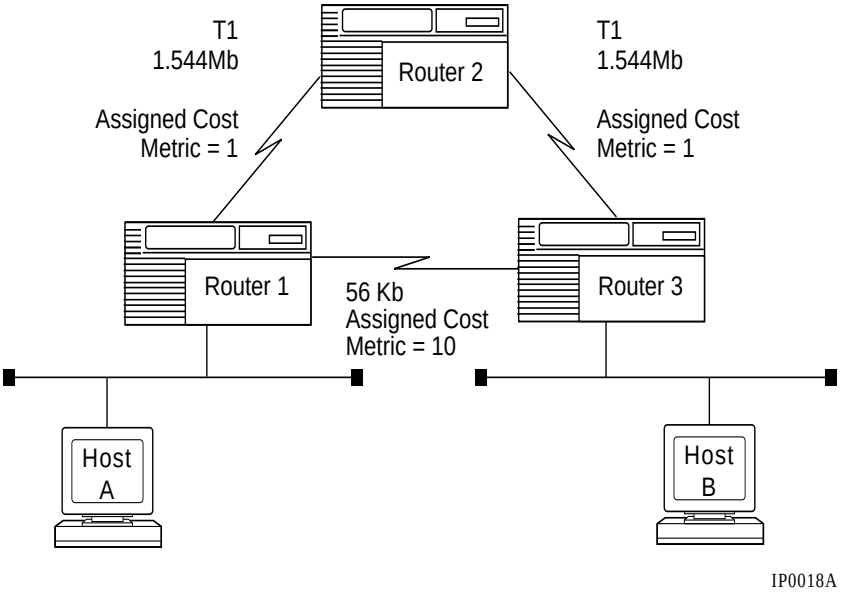

Figure 5-3. Configurable Cost Metrics Usage Example

There is an optimum cost for each type of network. We suggest the following values for this parameter:

Network Type//Bit Rate	Suggested Metric Cost
> = 100 Mb/s	1 (default)
Ethernet/802.3	10
E1	48
T1	65
64 Kb/s	1562
56 Kb/s	1785
19.2 Kb/s	5208
9.6 Kb/s	10416

Specifying the MTU Size

OSPF recognizes a maximum transmission unit (MTU) size for updates transmitted on an interface.

By default, when you configure OSPF on an interface, OSPF uses the MTU size specified for the type of network to which the interface is connected.

Using Site Manager, you can configure OSPF to

- Send packets no larger than the IP MTU size for Ethernet (1500).
- Use the MTU size you specify. The number you enter must be less than the IP MTU size for that physical interface.

Note: When running OSPF over a synchronous/PPP link, set the MTU size to a value less than the synchronous MTU size (1200). This allows all OSPF routes to be learned over the link.

Site Manager: MTU Size parameter: page A-83

Configuring OSPF on a Circuitless Interface

In some configurations, you may want the router's circuitless interface to be treated as the interface for an OSPF internal route. To do this, you can add OSPF to the circuitless interface. For information on the circuitless interface, see "Using the Circuitless IP Interface" on page 2-15.

Configuring a Neighbor on an NBMA Interface

In a nonbroadcast multiaccess network, neighbors are not learned dynamically. You can use Site Manager to enable and disable the neighbor configuration, supply the IP address for each neighbor, and specify the neighbor's priority.

Site Manager: Neighbor's IP Address parameter: page A-84

Configuring a Boundary Router

An OSPF boundary router

1. Receives information about routes outside the OSPF AS (using BGP, RIP, or another roouting protocol)
2. Formats this information in AS external (ASE) advertisements
3. Propagates the ASEs into the OSPF domain (using OSPF)

Enabling the Boundary Function

By default, the boundary function is disabled on the router. You can use Site Manager to configure the router as an OSPF boundary router.

Site Manager: AS Boundary Router parameter: page A-74

Configuring AS External Advertisements

Each ASE advertisement that a boundary router injects into the AS includes a Type 1 or Type 2 metric. The Type 1 metric is equivalent to the metric of the non-OSPF route. The Type 2 metric is either the metric of the non-OSPF route or the weight value calculated for that route.

Figure 5-4, for example, shows three routers in an OSPF domain. Router A and router B are both configured to generate ASE routes using the route weight as the Type 2 metric.

1. Boundary Router A learns a route to destination X via EGP.
2. Boundary Router A advertises the route to the internal router as an OSPF ASE route. The Type 2 metric in the advertisement contains the route weight value calculated for the EGP route to destination X.
3. Boundary Router B learns a route to destination X via BGP.
4. Boundary Router B advertises the route to the internal router as an OSPF ASE route. The Type 2 metric in the advertisement contains the route weight value calculated for a BGP route.
5. To determine the preferable route, the internal router compares the Type 2 metrics -- the EGP route weight and the BGP route weight.
6. The internal router selects the BGP route -- the route with the lower weight.

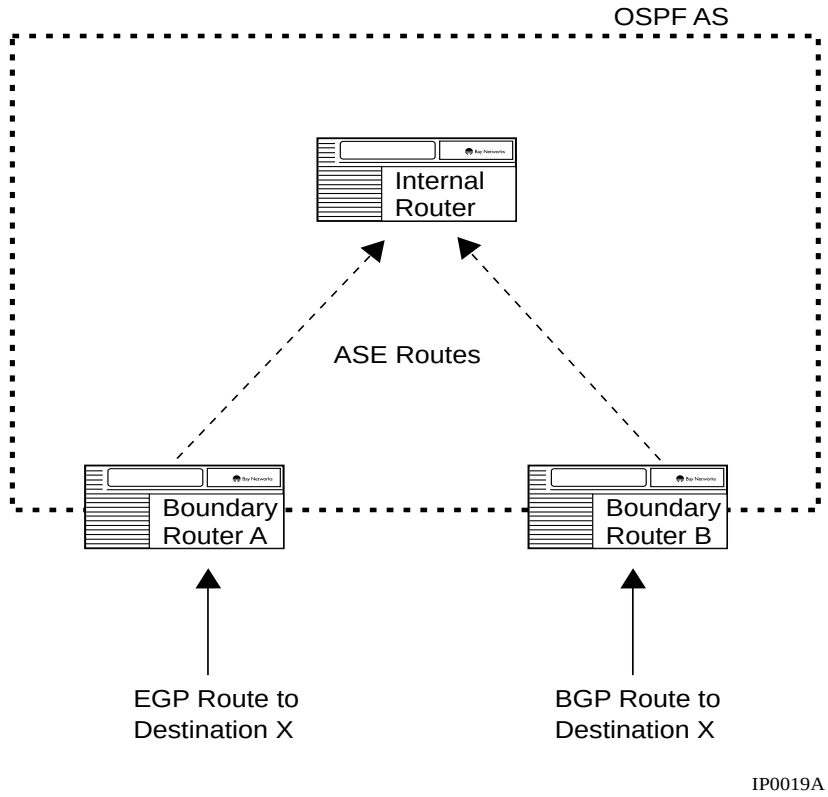

Figure 5-4. OSPF ASE Routes

By default, an OSPF boundary router generates a Type 2 metric for BGP, EGP, or RIP routes. For routes from all other sources, the boundary router generates a Type 1 metric.

Using Site Manager, you can configure a boundary router to use the route weight as the OSPF metric.

Site Manager: ASE Metric Support parameter: page A-75

Note: The route weight value will appear to be greater than the route's original metric. For this reason, all routers advertising a particular network must use the same metric type -- Type 1 or Type 2. If not, the router that receives the advertisements may choose the wrong route.

Using the External Route Tag in an ASE

An OSPF AS external route advertisement includes an external route tag field. This field allows boundary routers in an AS to exchange information about external routes. (The specific nature of this information is outside the scope of OSPF.)

By default, Bay Networks boundary routers that generate ASEs set the external route tag field to 0. For a boundary router running OSPF and BGP, the network administrator can configure OSPF to set the external route tag field with a value in accordance with RFC 1403, "OSPF/BGP Interaction."

In Figure 5-5, for example, boundary router A running OSPF and BGP learns external routes via BGP:

1. Router A generates an ASE describing the route. OSPF fills in the external route tag with BGP-specific information according to RFC 1403.
2. Router A injects the ASE into the AS, and OSPF routers flood the ASE throughout the AS.
3. Router B, which runs OSPF and BGP, receives the ASE. Router B generates a BGP update, using the contents of the external route tag to set the Origin and AS Path attributes in the update.

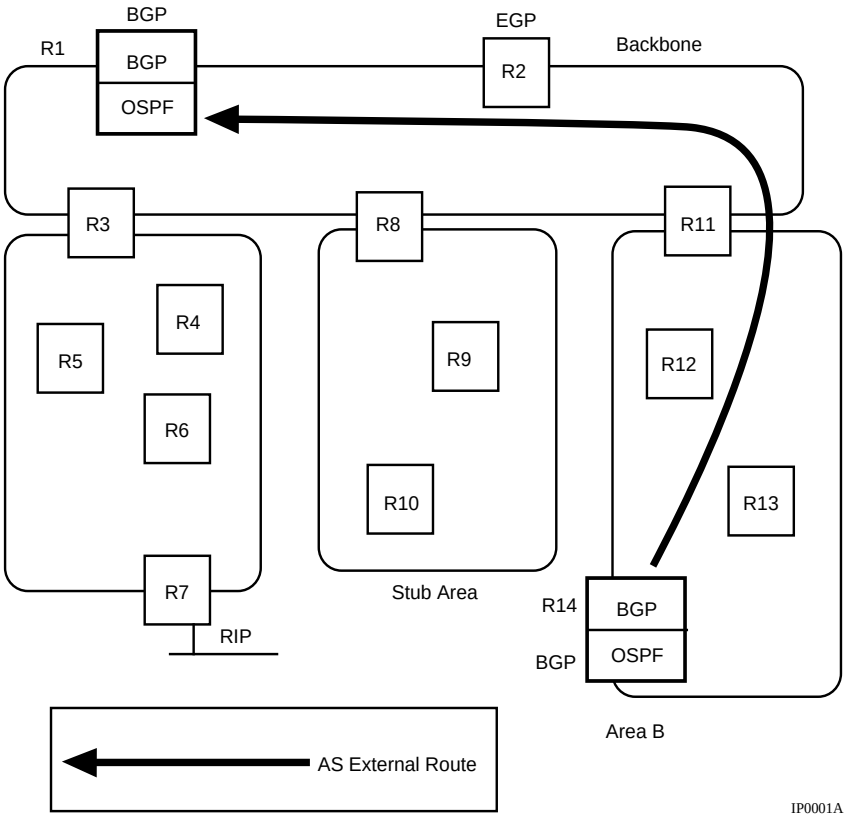

Figure 5-5. AS External Route Tag

You can use Site Manager to select the tag generation method.

Site Manager: Tag Generation Method parameter: page A-77

Configuring Authentication on an Interface

OSPF provides a measure of security for an area through the use of passwords. If an area is configured to use authentication, all OSPF interfaces configured in that area must be configured with a password. The password must be identical on each interface connected to the same network. Different networks can have different passwords.

In such an area, a router that receives a packet verifies the password before doing anything else with the packet. Unauthorized routers are not allowed to communicate with the OSPF system.

By default, authentication is disabled on each interface. You can use Site Manager to enable authentication on the interface.

If you have enabled authentication on the interface, you can use Site Manager to specify a password.

Site Manager: Authentication Type parameter: page A-86
Site Manager: Password parameter: page A-83

Configuring an Area Border Router

Each area is connected to the backbone by one or more border routers. In Figure 5-1, for example, R3 is a border router connecting Area 0.0.0.1 to the backbone. A border router can have connections to multiple areas. In Figure 5-6, for example, Border router 2 has an interface to a network in Area B and an interface to a network in Area C.

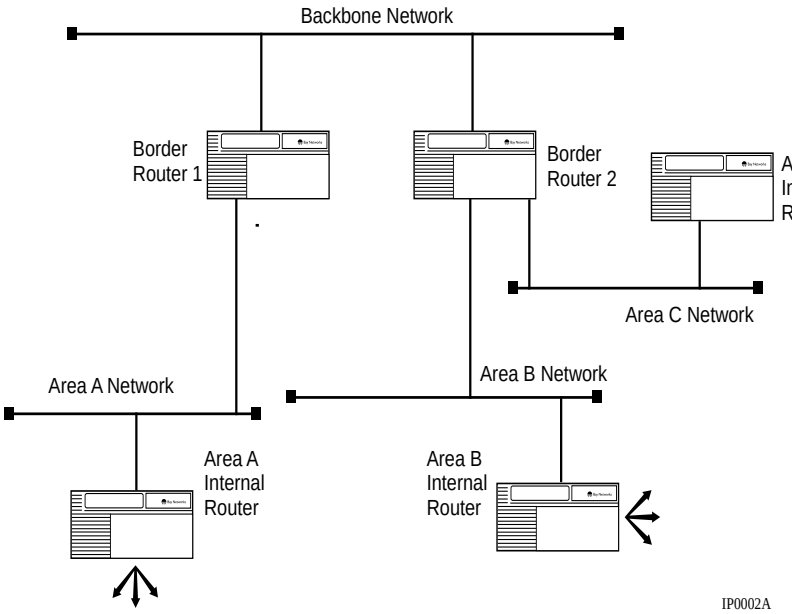

Figure 5-6. Area Border Router

Each border router in an AS

- Receives routing information from its attached areas, creates summaries of this information, and forwards the summaries to the backbone and to any other attached area. In Figure 5-6, for example, Border router 2 floods summaries from Area B to the backbone and Area C. Through the backbone, the summaries are forwarded to all other areas in the AS.
- Receives (via the backbone) summaries from other border routers, uses this information to create new routing summaries (which add in the cost of the backbone routes), and forwards the new summaries to its attached areas.

By definition, a border router has an interface to the backbone and interfaces to one or more other areas. To configure OSPF as a border router:

1. **Assign the backone ID (0.0.0.0) to an OSPF interface**
2. **Assign an area ID to another OSPF interface**

The features and functions of OSPF border routers are described in the following sections:

- Configuring a Virtual Backbone Link through a Transit Area -- page 5-32
- Defining a Range for a Summary Route -- page 5-34
- Configuring Border Services for a Stub -- page 5-34

Configuring a Virtual Backbone Link through a Transit Area

Every border router must have a connection to the backbone. This connection can be physical or virtual.

If the border router has an interface to a backbone network, that router is considered to be physically connected to the backbone. In Figure 5-6, Border router 1 and Border router 2 are both physically connected to the backbone.

In some cases, it may not be possible to configure a border router with an interface to a backbone network. If the router has an OSPF neighbor that is physically connected to the backbone, the router can use that neighbor to establish a *virtual link* to the backbone.

In Figure 5-7, for example, Border router 1 has lost its interface to the backbone network. In its place, the network administrator has configured an interface to a network in Area B. Through this network, Border router 1 now has a neighbor -- Border Router 2 -- that is connected physically to the backbone. The network administrator can use Border router 2 to configure a virtual link between Border Router 1 and the backbone.

An area that provides a virtual link between a border router and the backbone is considered to be a *transit area*. In Figure 5-7, Area B functions as a transit area.

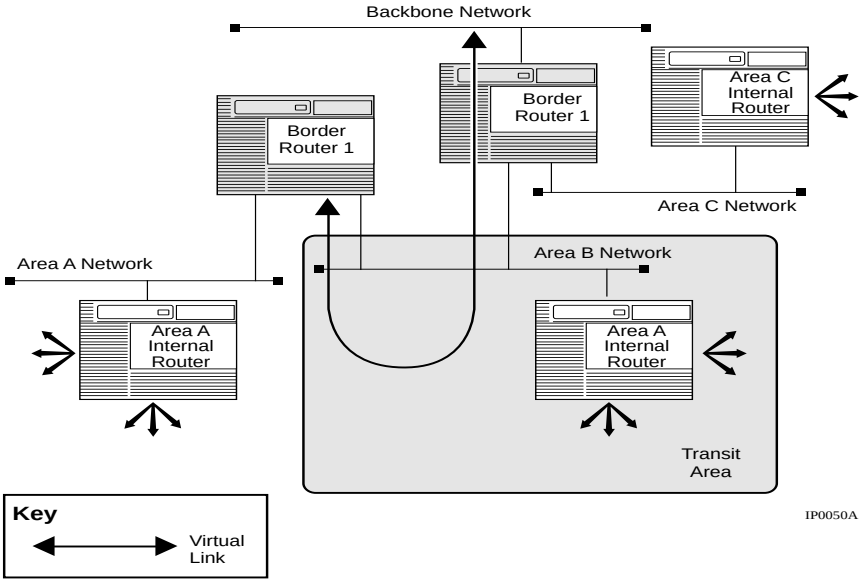

Figure 5-7. Virtual Link and Transit Area

To configure an interface to support a virtual link:

1. **Identify the transit area that supports the virtual link.**
2. **Identify the interface of the OSPF neighbor at the other end of the virtual link.**

Once you have defined the virtual link, you can

- Enable and disable the virtual link.
- Specify a transit delay, a retransmit interval, a Hello interval, and a dead interval for the link.
- Specify a password.

You can use Site Manager to perform these operations on an OSPF virtual link.

Site Manager: Transit Area ID parameter: page A-90
Site Manager: Neighbor's Router ID parameter: page A-90
Site Manager: Enable parameter: page A-91
Site Manager: Transit Delay parameter: page A-91
Site Manager: Retransmit Interval parameter: page A-92
Site Manager: Hello Interval parameter: page A-92
Site Manager: Dead Interval parameter: page A-93
Site Manager: Password parameter: page A-93

Defining a Range for a Summary Route

Border routers generate summary advertisements for their attached areas. Each summary advertisement specifies a range of destinations in an area. An area range specification consists of a network address and a variable-length mask. For example, a summary advertisement for the destination 140.191.0.0 with a mask of 255.255.0.0 describes a single route to the collection of destinations 140.191.0.0 to 140.191.255.255. When a packet is forwarded, it is always forwarded to the network that is the best (longest or most specific) match for the packet's destination.

An address range is associated with an interface to an area. You can use Site Manager to define a range for an area.

Site Manager: Range Net parameter: page A-88
Site Manager: Range Mask parameter: page A-88

Configuring Border Services for a Stub

A stub area does not import ASEs and may or may not import internal route summaries.

In place of routes to destinations outside the stub, a border router connected to a stub injects a default route advertisement. When an internal router encounters a datagram addressed to a destination outside the stub, the router forwards it to the border router specified in the default route advertisement.

Assume, for example, that the stub area in Figure 5-1 has been configured to import no internal or external routing information. Border router 8 receives ASEs and internal summaries from its interface to the backbone. However, Border router 8 does not forward the ASEs or summaries to the stub. Instead, it injects a default route that internal routers use to forward datagrams to destinations beyond the stub.

By default, a border router imports ASEs into its attached areas. You can use Site Manager to disable and re-enable the importing of ASEs for a stub area.

By default, a border router that injects a default route into a stub area assigns a cost metric of 1 to that default route. You can use Site Manager to specify a cost metric (overriding the default).

By default, a border router injects network summaries into an attached stub area. You can use Site Manager to disable the injection of summaries.

Site Manager: Import AS Extern parameter: page A-86
Site Manager: Stub Metric parameter: page A-87
Site Manager: Import Summaries parameter: page A-87

Applying IP Policies to OSPF Interfaces

According to the OSPF standard, all routers in a given area must maintain a similar routing database. To ensure the integrity of the database, the network administrator must not use IP policies to manipulate link-state advertisements before propagating them on an interface.

There are two ways, however, in which an IP policy can be applied to an OSPF interface:

- An OSPF announce policy can be used on a boundary router to control which self-originated external routing updates are placed in the link-state database for distribution according to the OSPF standard. The announce policy affects what other routers learn but only with regard to the local boundary router's self-originated information.

- An OSPF accept policy can be used on a router to control which OSPF non-self-originated external routing information is stored in the routing table. This accept policy controls only what the local router uses; it does not affect the propagation of non-self-originated external information to other routers.

Chapter 6
Customizing BGP Services

This chapter consists of the following sections describing the Bay Networks implementation of the Border Gateway Protocol (BGP).

- BGP Concepts and Terminology -- page 6-2
  - Interior BGP in Intra-AS Routing -- page 6-3
  - IBGP in a Transit AS -- page 6-3
  - BGP Path Attributes -- page 6-5
  - BGP-4 Local Preference Values -- page 6-7
  - BGP Implementation Notes -- page 6-8
- Configuring BGP on the Router -- page 6-9
  - Supplying a BGP Identifier -- page 6-9
  - Configuring BGP for Intra-AS Routing -- page 6-10
  - Setting the Interval Timer -- page 6-11
  - Allowing Redundant Connections -- page 6-11
  - Enabling Multihop Connections -- page 6-12
  - Disabling Dynamic Policy Configuration -- page 6-12
  - Configuring BGP as a Soloist -- page 6-13
- Establishing a Peer-to-Peer Connection -- page 6-13
  - Defining the Peer-to-Peer Connection -- page 6-13
  - Initiating a Peer-to-Peer Connection -- page 6-16
  - Negotiating the BGP Version -- page 6-17
- Designing a Route Selection Strategy -- page 6-21
  - Assigning AS Weight and Class Values -- page 6-21
  - Configuring Routing Policies -- page 6-22
- OSPF/BGP Interaction -- page 6-25
- Configuring BGP Message Logging -- page 6-25

BGP Concepts and Terminology

BGP is an exterior gateway protocol designed to exchange network reachability information with other BGP systems in other autonomous systems.

BGP exchanges routing information in the form of routing updates. An update includes a network number, a list of autonomous systems that the routing information has passed through (the AS path), and a list of other path attributes.

Figure 6-1 shows two autonomous systems: AS1 and AS2. Networks within AS1 and AS2 are connected by routers running an interior gateway protocol -- in this case, OSPF. AS1 and AS2 are connected by routers that run an exterior gateway protocol -- BGP -- in addition to OSPF.

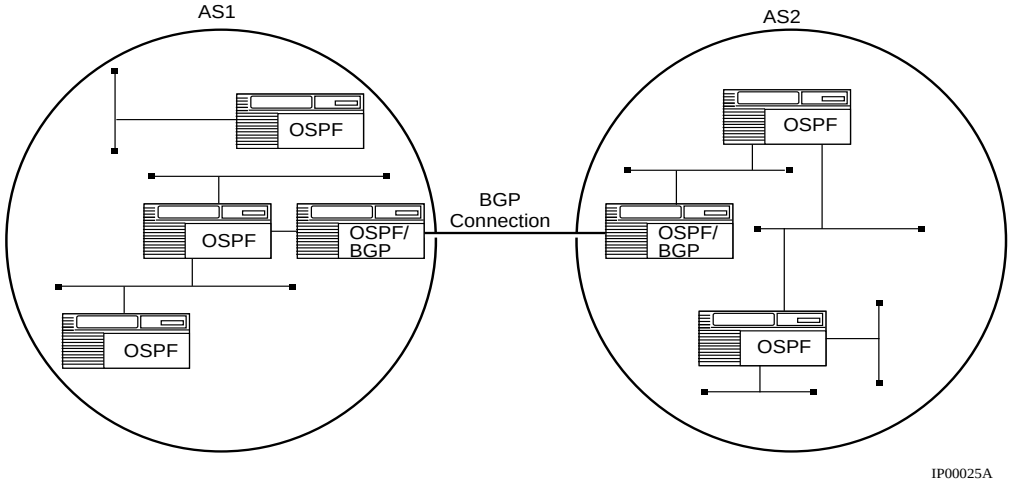

Figure 6-1. BGP Connecting Autonomous Systems Running OSPF

A BGP router employs a BGP speaker, which is an entity within the router that transmits and receives BGP messages and acts upon them. A BGP speaker forms a neighbor relationship with another BGP speaker by establishing a *peer-to-peer session* over a common inter-AS network.

An autonomous system can include one or more BGP speakers that provide external route information for the networks within the AS. An AS containing a single BGP speaker with a single external BGP connection is a stub AS. The BGP speaker is providing external route information for the networks contained within its AS only.

Bay Networks supports two versions of the Border Gateway Protocol: BGP-3 and BGP-4.

- BGP-3 assumes that each advertised network is a natural class network (A, B, or C) based on its high-order bits. BGP-3 cannot advertise subnets or supernets.
- BGP-4 has no concept of address classes. Each network listed in the Network Layer Reachability Information (NLRI) portion of an Update message contains a prefix length field, which describes the length of the mask associated with the network. This allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless interdomain routing (CIDR) deployment possible.

Interior BGP in Intra-AS Routing

Bay Networks implements Interior BGP (IBGP) intra-AS routing. Under IBGP, each router in the AS runs an IGP for internal routing updates and also maintains an IBGP connection to each BGP border router. The IBGP information is used in conjunction with the IGP route to the authoring BGP border router to determine the next hop to use for external networks.

No BGP information is carried by the IGP. Each router uses IBGP exclusively to determine reachability to external networks. When an IBGP update for a network is received, it can be passed on to IP for inclusion in the forwarding tables only if a viable IGP route to the correct border gateway is available.

IBGP in a Transit AS

An AS with more than one BGP speaker can use IBGP to provide a transit service for networks outside the AS. An AS that provides such a service for BGP speakers is known as a *transit AS* (see Figure 6-2).

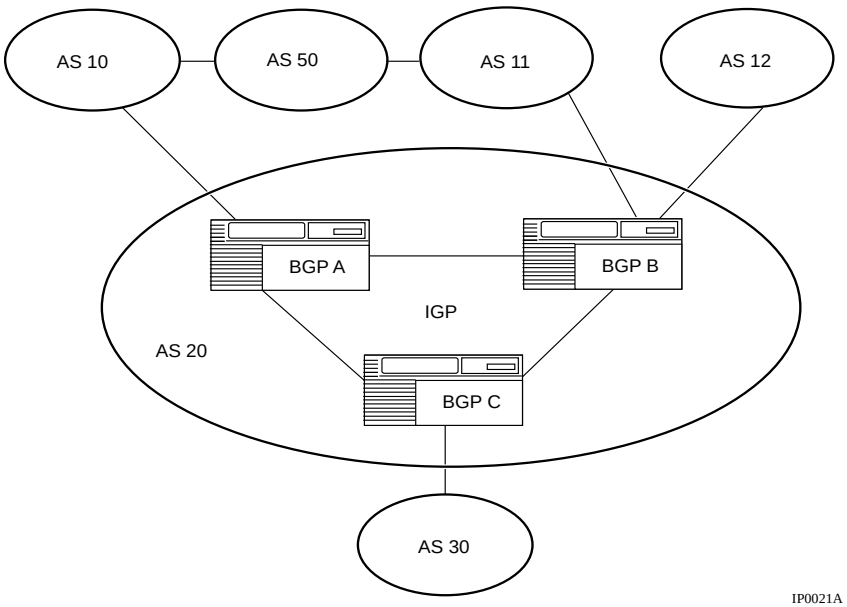

Figure 6-2. Transit Autonomous System

In Figure 6-2, Autonomous System 20 is the transit AS. It is providing information about its internal networks, as well as transit networks, to the remaining ASs. The IBGP connections between BGP routers A, B, and C are necessary to provide consistent information to the ASs.

BGP Path Attributes

A BGP-3 Update message includes a variable-length sequence of path attributes. Each attribute entry consists of an attribute value and a field describing the attribute. Table 6-1 lists the mandatory and optional BGP-3 path attributes.

Table 6-1. BGP-3 Path Attributes

Attribute	Description
AS Path	Mandatory attribute containing a list of the ASs that must be traversed to reach the given destinations.
Origin	Mandatory attribute containing one of the following values: IGP (the path is valid all the way to the IGP of the originating AS), EGP (the path was advertised using EGP by the last AS in the AS path), or Incomplete (the path is valid only to the last AS in the AS path).
Next Hop	Mandatory attribute that defines the IP address of the router to use as a next hop for the advertised destinations.
Inter-AS	Optional attribute used to choose between paths to the destinations listed.
Unreachable	Discretionary attribute used to indicate destinations that have become unreachable.

The BGP-4 update message has the same format and contains the same mandatory attributes as the BGP-3 update message with the following additions:

- In place of the Unreachable attribute that BGP-3 includes as part of the path attribute description, the BGP-4 update includes an Unreachable field. This field specifies destinations that have become unreachable.
- In place of the BGP-3 optional attributes, a BGP-4 update message can include the optional attributes described in Table 6-2.

Table 6-2. BGP-4 Optional Path Attributes

Attribute	Description
Multi-Exit Discriminator	Optional attribute used to choose between paths to the destinations listed.
Local Preference	Optional attribute allowing AS border routers to indicate the preference they have assigned to a chosen route when advertising it to IBGP peers.
Atomic Aggregate	Optional attribute used to ensure that certain network layer reachability information (NLRI) is not deaggregated.
Aggregator	Optional attribute identifying which AS performed the most recent route aggregation. The attribute contains the last AS number that formed the aggregate route followed by the IP address of the BGP speaker that formed the aggregate route.
Route Clusters	Optional attribute listing the route clusters that may be traversed to reach a given destination.
Advertiser	Optional attribute identifying which border router injected the route.
BGP Community	Optional transitive attribute identifying the communities to which the route belongs. (A community is a group of destinations that share some common property.)

BGP-4 Local Preference Values

BGP-4 update messages include a Local Preference attribute that allows an AS border router to assign a preference value to a route when advertising it to IBGP peers. The calculation of the Local Preference attribute is implementation-specific. A higher value indicates that the route is more preferred.

The router uses the following equations to calculate a value for the Local Preference attribute:

local preference = 8191 – *origin value* – *AS path weight*

where *origin value* is 0 for routes with an Origin Path attribute of IGP and 4096 otherwise, and *AS path weight* is a sum of weight values associated with AS numbers listed in the route's AS Path attribute. These weight values can be configured and default to 8.

A steep penalty is applied to routes that are advertised with an Origin attribute other than IGP -- that is, EGP or Incomplete.

For an OSPF internal route or a direct route, the Local Preference attribute is set to

local preference = (8191 + 256 – (*metric* & 255))

where *metric* is the OSPF metric for an OSPF route or the configured cost for a direct route.

For a RIP route, an EGP route, an OSPF ASE route, or a static route, the Local Preference attribute is set to

local preference = (256 – *metric*)

where *metric* is the RIP metric for a RIP route, the EGP metric for an EGP route, the OSPF metric for an OSPF ASE route, or the configured cost for a static route.

Note that Local Preference values for OSPF internal routes and direct routes are higher than the Local Preference values calculated for BGP routes.

BGP Implementation Notes

This section provides you with some guidelines that you should follow when you configure BGP. If you do not follow these guidelines, BGP will either not work efficiently or will become disabled on the interfaces involved.

- BGP will not operate with an IP router in nonforwarding (host-only) mode. Make sure that the routers you want BGP to operate with are in forwarding mode.
- If you are using BGP for a multihomed AS (one that contains more than one exit point), we strongly encourage you to use OSPF for your IGP and BGP for your sole exterior gateway protocol, or use intra-AS IBGP routing.

  If OSPF is the IGP, you should also use the default OSPF tag construction. Using EGP or modifying the OSPF tags makes network administration and proper construction of BGP path attributes more difficult.
- For any router supporting both BGP and OSPF, the OSPF router ID and the BGP identifier must be the same.

Configuring BGP on the Router

The following sections describe settings that affect the way BGP runs on the router.

Enabling and Disabling BGP

When you start BGP on the router, BGP is automatically enabled for both BGP-3 and BGP-4 peer-to-peer connections.

You can use Site Manager to disable and re-enable BGP on the router. You can also use Site Manager to disable and re-enable BGP-3 and BGP-4.

Site Manager: BGP Enable parameter: page A-2
Site Manager: Enable parameter: page A-5 (BGP-3)
Site Manager: Enable parameter: page A-6 (BGP-4)

Supplying a BGP Identifier

The BGP identifier is the IP address of an interface on this router. There is no default for this parameter. You must supply a BGP ID, using the IP address of one of the router's IP interfaces.

You can use Site Manager to supply a BGP identifier for the router.

Site Manager: BGP Identifier parameter: page A-2

Identifying the Local AS

Each autonomous system in the internet has a unique AS ID.You can use Site Manager to supply the ID of the AS in which the BGP router is located.

Note: If both BGP and OSPF are running on the router, then the OSPF router ID must be equivalent to one of the configured IP interfaces.

Site Manager: BGP Local AS parameter: page A-2

Configuring BGP for Intra-AS Routing

By default, BGP supports IBGP intra-AS sessions. (See “Interior BGP in Intra-AS Routing” on page 6-3.)

You can use Site Manager to disable and re-enable the feature.

Site Manager: BGP Intra-AS parameter: page A-3

If IBGP is enabled, you can use Site Manager to specify the types of routes that BGP advertises in IBGP sessions. By default, BGP propagates only routes learned from external BGP peers.

Site Manager: BGP From Protocols parameter: page A-3

A BGP transit AS should use IBGP intra-AS routing. A stub or multihomed AS usually does not use IBGP routing.

Setting the Interval Timer

BGP injects external BGP routes into the routing table. The default minimum interval between injections is five seconds.

You can use Site Manager to specify the minimum number of seconds between injections.

Site Manager: BGP Interval Timer parameter: page A-3

Allowing Redundant Connections

By default, BGP can maintain only one TCP session with a remote BGP peer. If the remote peer attempts to establish another session on another physical connection, BGP rejects the session. (BGP uses a collision-detection method based on router ID to check for redundant sessions.)

You can use Site Manager to allow TCP sessions on multiple physical connections (that is, to disable redundancy checking).

The advantage of a peer-to-peer configuration with multiple sessions on multiple physical connections is redundancy -- if one connection fails, the peers can communicate over another link. The disadvantage is that such a configuration results in multiple copies of each route.

Site Manager: BGP Collision Detect parameter: page A-4

Enabling Multihop Connections

By default, BGP enforces the one-hop rule for BGP peers -- the remote peer must be located on a directly attached network.

You can use Site Manager to override the restriction and allow multihop connections.

Caution: Enabling multihop BGP connections is dangerous because it can cause BGP speakers to establish a BGP connection that traverses a third-party AS, which may violate policy considerations and may also introduce forwarding loops.

Site Manager: Multi-hop EBGP Connection parameter: page A-4

Disabling Dynamic Policy Configuration

By default, BGP configures IP policies dynamically. This means that if you modify a policy, BGP dynamically re-evaluates all affected routes in the light of the modified policy. BGP then sends the appropriate withdraw or update to the affected peers. BG P maintains records of which routes have been sent to which peer. This allows for precise determination of which routes must be sent and which must be withdrawn.

You can use Site Manager to disable and re-enable dynamic policy configuration.

Site Manager: Dynamic Policy Change parameter: page A-5

If you modify an IP policy with this feature disabled, BGP restarts all BGP connections. There is no advantage to disabling dynamic policy configuration. Disabling this parameter will significantly impact BGP protocol operation overhead and network stability.

Configuring BGP as a Soloist

By default, BGP runs on every slot on the router. If you have configured ISP mode, BGP runs as a soloist. The soloist will run on one slot, as determined from the BGP soloist slot mask. We recommend that the slot mask include only nonforwarding slots, so that BGP operations (route calculation, for example) occur on one slot while the other slots maintain maximum forwarding capability.

Site Manager: ISP Mode parameter: page A-45

In soloist mode, if the slot on which the soloist is running fails, BGP runs on an eligible slot. By default, BGP considers all slots with IP interfaces to be eligible slots. You can use Site Manager to specify the eligible slots.

Site Manager: Soloist Slots parameter: page A-5

Establishing a Peer-to-Peer Connection

A BGP speaker forms neighbor relationships with other BGP speakers. This happens when a BGP speaker establishes a TCP connection to a BGP peer (which is simply the BGP speaker at the other end of the connection), based on local configuration information.

Defining the Peer-to-Peer Connection

To define a peer-to-peer connection, you specify

- The address of the local IP interface
- The address of the remote IP interface. If the peer is in a remote AS, this address must be on the same subnet as the local address.
- The AS number of the autonomous system in which the remote BGP peer is located

You can use Site Manager to supply this information.

Site Manager: Local Address parameter: page A-7
Site Manager: Peer Address parameter: page A-6
Site Manager: Peer AS parameter: page A-6

Using the Circuitless IP Interface for a Peer-to-Peer Connection

In configuring a peer-to-peer connection for BGP speakers, you specify a local peer address (the address of a local IP interface) and a remote peer address (the address of a remote IP interface).

In situations where BGP speakers reside on routers that have multiple network connections over multiple IP interfaces (the typical case for IBGP speakers), consider using the address of the router's circuitless IP interface as the local peer address.

By using the address of the circuitless IP interface as the local peer address in an IBGP configuration, you are ensuring that BGP is reachable as long as there is an active circuit on the router.

Configuring BGP Peers over an Unnumbered Point-to-Point Link

A BGP peer-to-peer connection cannot be configured directly on an unnumbered interface. To establish a connection, each side of the connection must be associated with a numbered interface.

For example, consider the two routers in Figure 6-3. Router A and Router B are connected by a point-to-point network using unnumbered interfaces. Both routers are configured with BGP.

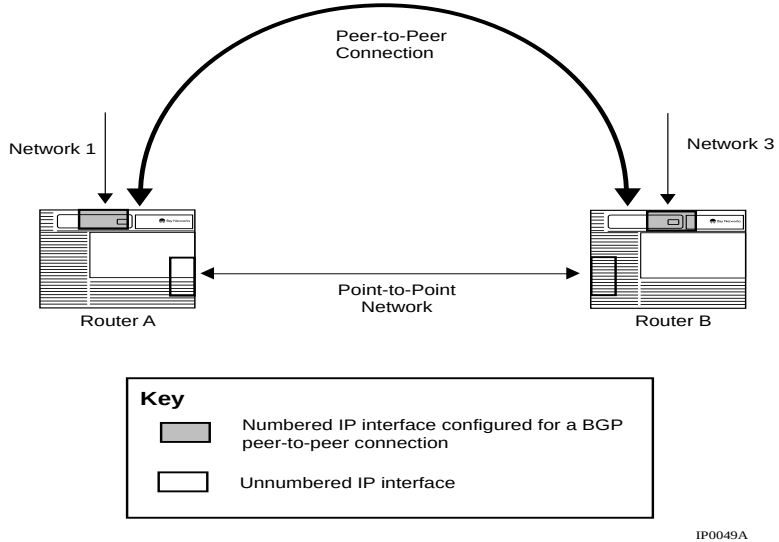

Figure 6-3. BGP over an Unnumbered Point-to-point Link

To establish a peer-to-peer connection between Router A and Router B, you must

1. **Choose a numbered interface on each router for the peer-to-peer connection. Note the network/subnet that each interface is on.**

   In Figure 6-3, Router A has a numbered interface to Network 1. Router B has a numbered interface to Network 3. The network administrator has chosen these two interfaces to support the peer-to-peer connection.

2. **If the two routers are in different ASs, enable multihop EBGP connections.**

3. **If no IGP protocol (RIP or OSPF) is running over the unnumbered link, configure a static route on each router to the other router's network and subnet.**

   Since the routers do not share a numbered subnet, each BGP peer needs to know a route to the network/subnet of the interface that the other BGP peer is using. If there is an IGP protocol (RIP or OSPF) running over the unnumbered link, RIP or OSPF will learn the route and store it in the routing table. Otherwise, you need to configure a static route on each router to the other BGP peer's network/subnet. The route should point to the unnumbered link.

4. **Configure the BGP connection on each router.**

Initiating a Peer-to-Peer Connection

A BGP speaker that wants to initiate peer-to-peer connections periodically issues an Open message.

BGP speakers respond to connection requests by returning an Open message. In Figure 6-4, for example, BGP Speaker A sends an Open message to BGP Speaker B to request a connection; BGP Speaker B responds by sending an Open message to BGP Speaker A.

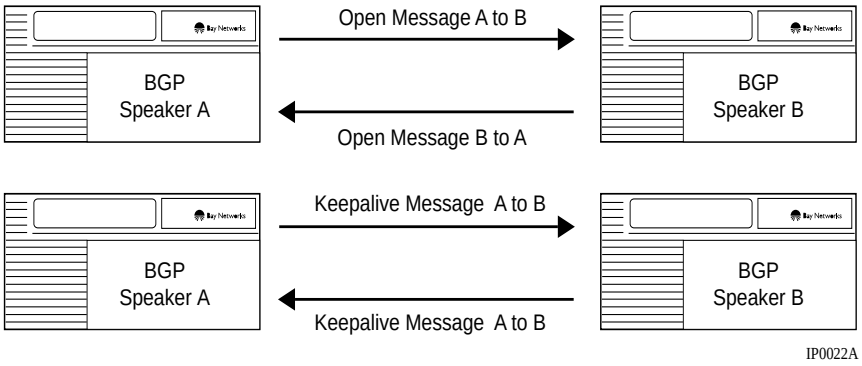

Figure 6-4. Establishing and Confirming a Connection between BGP Peers

All BGP speakers respond to connection requests from other speakers.

By default, BGP attempts to initiate a connection on each interface configured for peer-to-peer communications. If the attempt is unsuccessful, BGP retries every 120 seconds.

You can use Site Manager to specify a retry interval or disable the initiation function.

Site Manager: Connect Retry Timer parameter: page A-9

Negotiating the BGP Version

BGP peers negotiate the version of BGP that they will use to exchange routing information. If you enable both BGP-3 and BGP-4, the router first attempts to use BGP-4. If the BGP peer is not a BGP-4 speaker, the router uses BGP-3.

By default, BGP considers BGP-4 as both the minimum and maximum acceptable version for negotiation.

You can use Site Manager to specify BGP-3 as the minimum or maximum acceptable version.

Site Manager: Min BGP Version parameter: page A-7
Site Manager: Max BGP Version parameter: page A-8

Keeping the Connection Alive

Once a session has been established, BGP peers periodically issue a Keepalive message to maintain the connection. By default, BGP issues a Keepalive message every 30 seconds.

You can use Site Manager to specify how often BGP issues a Keepalive message on this peer connection or to disable the Keepalive function.

Site Manager: Keepalive Timer parameter: page A-10

Specifying an Update Interval for the Connection

Once a connection is established, the BGP speaker uses one or more Update messages to send the entire IP routing table (compliant with local BGP export policies). BGP, however, does *not* require the entire routing table to be sent again. Therefore, the BGP speaker must keep a current version of the routing information received from of all of its peers for as long as the connection to each peer is valid. This information will be updated via Update messages whenever changes occur.

By default, BGP examines the routing table for changes every 5 seconds. If a change has occurred, BGP issues an Update message on the connection.

You can use Site Manager to specify a value for the external advertisement timer.

Site Manager: External Advertisement Timer parameter: page A-8

Specifying a Holddown Time

The holddown time is the amount of time either peer will wait for a Keepalive or Update message before declaring the connection down.

A BGP speaker that is initiating a connection inserts a holddown time value into the Open message. The responding peer responds with an Open message that also contains a holddown time value. If the BGP speakers establish a session, they use the lesser value (which must be greater than 2). There are two exceptions to this rule:

If one peer sends a zero holddown time, the peers use the nonzero holdtime on the session.

If both peers send zero holddown times, the peers observe no holddown time on the session.

By default, BGP inserts a value of 90 seconds into the Open message. You can use Site Manager to specify a holddown time value or disable the holddown function.

Site Manager: Holdtime parameter: page A-9

Setting a Minimum AS Origination Interval

By default, a BGP speaker that issues an update to advertise a change in the AS must wait at least 15 seconds before advertising a subsequent change. You can use Site Manager to specify a different interval.

Site Manager: Min AS Origination Interval parameter: page A-10

Overriding the Local AS Number

By default, a BGP speaker that issues an open message to initiate a peer-to-peer session uses the AS number that you set with the Local AS parameter.

You can use Site Manager to include a different AS number (overriding the default) or use the AS number you specified in the Local AS parameter.

Site Manager: Local AS to Advertise to Peer parameter: page A-10

Specifying a Maximum Update Size

By default, a BGP speaker sends updates with a maximum size of 800 bytes.

You can use Site Manager to specify a maximum update size (overriding the default).

Note that, if the update message that is used to advertise a single route is larger than the configured message size, the actual message size can exceed the configured value.

Site Manager: Peer Max Update Size parameter: page A-11

Setting the Route Echo Switch

The peer route echo switch controls the way the router echoes a BGP route that is selected for forwarding. (Echoing in this case means advertising the route back to the peer from which it was received.)

By default, the router advertises the route back as reachable and includes the local AS.

You can use Site Manager to configure BGP to echo the route as UNREACHABLE/withdrawn.

Site Manager: Peer Route Echo Switch parameter: page A-11

Designing a Route Selection Strategy

A BGP speaker must, at times, evaluate and compare different paths to a destination network to determine the best path. Because all border routers must provide the same view of the AS to external ASs, having a selection strategy that is consistent in the router, and that can be consistent across all border routers, is very important in BGP. To select the best available path, BGP uses AS weights and classes and IP policies.

Assigning AS Weight and Class Values

You can assign a weight class to any AS number and a weight value to a weight class.

AS weight classes allow a network administrator to assign multiple weight values to the same AS. This feature allows the administrator to consider an AS path differently for different networks. For example, consider a situation in which two networks -- 192.32.1.0 and 192.32.2.0 -- are both reachable by two paths. The first path to each network shares a common AS -- AS 5. The second path to each network also shares a common AS -- AS 10. If the administrator for some reason wants to favor AS 5 in the path to 192.32.1.0 and AS 10 in the path to 192.32.2.0, she can assign one weight class to the AS as a component in the path to network 192.32.1.0 and another class to the AS component in the path to 192.32.2.0.

An assigned weight can range from 1 to 15 plus an infinity value. Weights provide a way either to prefer or to avoid routes that pass through certain ASs. The weights of each AS in a path are added, and the path with the smallest total weight is the preferred path. Any path containing an AS weight of infinity will be avoided.

When a BGP router receives a new route, it is evaluated against any existing accept policies. If after this evaluation, the path still is to be used, the total weight of the path is calculated.

AS weights should be configured the same on all BGP routers in an AS.

You can use Site Manager to assign a weight and a weight class to an AS.

Site Manager: Weight Value 1 parameter: page A-12
Site Manager: Weight Value 8 parameter: page A-14

Configuring Routing Policies

BGP accept and announce policies govern which routes a router uses and which routes it propagates to other routers.

Note: By default, an external BGP-3 or BGP-4 speaker will neither advertise any routes to a peer, nor inject any routes into its IGP. Route policies must be configured to enable any route advertisement.

So that every BGP border router within an AS comes to the same decision in constructing path attributes for an external path, route policies must be coordinated between all of the BGP speakers within an AS. It is suggested that the accept and announce policies on all IBGP connections accept and propagate all routes. On external BGP connections, consistent routing policy decisions should be made.

Note: In addition to announce and accept policies, Bay Networks supports import and export filters for BGP-3. Import and export filters provide a subset of the parameters provided by the policies. In a future release, support for import and export filters will be discontinued.

When a BGP speaker receives a route in an Update message, it applies any local routing policies to determine whether the router will use the route and whether it will propagate the route to other routers. Then, if the route can be used, it is compared against routes from other protocols and possibly included in the forwarding table.

You set values in BGP-4 accept and announce policy parameters to match and, in some cases, override the attribute values contained in inbound and outbound update messages.

Configuring a BGP-4 AS Pattern-Matching Policy

In the Bay Networks implementation of AS path pattern matching for Version 10.0, the user must know regular expression syntax. For Version 11.0, we have provided a simpler syntax. This syntax is defined in Table 6-3.

IP software passes this simple form of regular expressions to a preprocessor that expands the simple expression into a generic regular expression.

There are two additional changes in AS pattern matching syntax for Version 11.0.

- AS pattern matching for Version 11.0 uses a new symbol to denote AS path segments of the type AS_SET. In Version 10.0, AS sets are denoted by the symbols '`. Beginning with Version 11.0, AS sets are denoted by {}.
- Version 10.0 requires that the user enter a space delimiter after entering the last AS in an AS pattern matching expression (for example, “555 $”. Beginning with Version 11.0, the extra space delimeter at the end of an AS list is no longer necessary (for example, “555$”).

Table 6-3. Characters in AS-Path Pattern Matching

Symbol or Operator	Meaning
<	Denotes beginning of an AS SEQUENCE segment.
>	Denotes end of an AS SEQUENCE segment.
{	Denotes beginning of an AS SET segment.
}	Denotes end of an AS SET segment.
<*seq*>{*set*}	denotes an AS path containing a sequence in the first segment and a set in the second segment.
^	Denotes the following pattern occurs at the beginning of the AS path.
$	Denotes the preceding pattern occurs at the end of the AS path.
|	Logical OR - match this or that.
_*X*_	Match exactly AS = “*X*”
_*X*	Match AS pattern beginning with *X* (for example, ”_99” matches 99 991 9934)
*X*_	Match AS pattern ending with *X* (for example, “99_” matches 99 199 23299)

You can use simplified regular expressions to configure BGP-4 accept and announce policies for AS pattern matching.

To configure a BGP-4 accept policy with Site Manager:

1. **Select Protocols > IP > Policy Filters BGP-4 > Accept Policy**
2. **Click on the AS Pattern parameter.**
3. **Use the simplified regular expressions to specify the pattern matching filter.**

To configure a BGP-4 announce policy with Site Manager:

1. **Select Protocols > IP > Policy Filters BGP-4 > Announce Policy**
2. **Click on the AS Pattern parameter.**
3. **Use the simplified regular expressions to specify the pattern matching filter.**

Best Route Calculation for Equal Routes

BGP uses the following rules (tie breakers) to choose between two equal BGP routes:

1. Choose the route with the lower route weight.
2. Choose the route with the higher Local Preference attribute.
3. Choose the route with the lower Inter-AS Metric attribute (if both routes include this optional attribute).
4. Choose the route with the lower interior cost to the Next Hop.
5. Choose external BGP over IBGP.
6. Choose the route with the lower BGP identifier.
7. Choose the route with the lower BGP connection remote address.
8. Choose the route with the lower BGP connection local address.

OSPF/BGP Interaction

RFC 1403 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. For routers running both protocols, the OSPF router ID and the BGP identifier must be an IP address and must be identical. A route policy must be configured to allow BGP advertisement of OSPF routes.

Interaction between BGP-4 and OSPF includes the ability to advertise supernets to support classless interdomain routing (CIDR). BGP-4 allows interdomain supernet advertisements. OSPF can carry supernet advertisements within a routing domain.

Configuring BGP Message Logging

Site Manager allows you to control the event messages that BGP sends to the log file by specifying

- Local and remote address of a peer-to-peer session or sessions
- Message severity level: fault, warning, information, trace, or debug, or all levels
- BGP message type: Open, Update, Notification, or Keepalive

Use BGP message logging parameters to limit the volume of debug-level messages that BGP generates and logs. If you allow BGP to log all debug-level events, the messages that BGP generates will quickly overrun and overwrite the log file.

Chapter 7
Customizing EGP Services

This chapter, which explains how to configure the Exterior Gateway Protocol, contains the following sections:

- EGP Concepts and Terminology -- page 7-1
- Configuring EGP on the Router -- page 7-6

  Enabling and Disabling EGP -- page 7-6
  Supplying a Local AS Number -- page 7-6
  Specifying the Neighbor's Address -- page 7-7
  Specifying the Gateway Mode -- page 7-7

- Configuring a Neighbor -- page 7-6

  Enabling and Disabling the Neighbor Relationship -- page 7-8
  Selecting the Acquisition Mode -- page 7-8
  Selecting the Poll Mode -- page 7-8
  Setting Neighbor Timers -- page 7-9

EGP Concepts and Terminology

EGP-2 is an exterior gateway protocol used to exchange network reachability information between routers in different autonomous systems. In each, AS routers share routing information using one or more interior gateway protocols -- for example, RIP or OSPF. The routers that serve as endpoints of a connection between two ASs run an exterior gateway protocol, such as EGP-2 (see Figure 7-1).

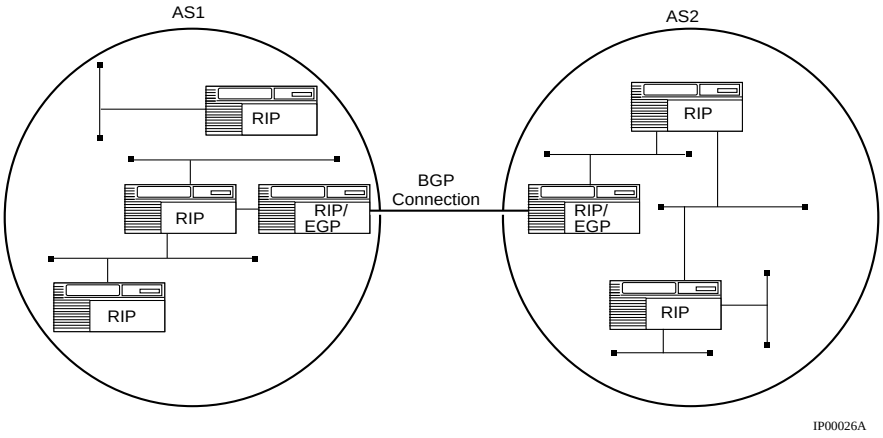

Figure 7-1. EGP Connection between Two Autonomous Systems Running RIP

The Bay Networks implementation of EGP complies with RFCs 827 and 904. It runs over the same LAN and WAN media/protocols that IP runs over, including Ethernet, Token Ring, Synchronous, Wellfleet Proprietary Synchronous, Frame Relay, SMDS, X.25 (DDN, PDN, Pt-to-Pt), ATM PVC, FDDI, T1, E1, HSSI, and PPP.

Note: EGP assumes that each advertised network is a natural class network (A, B, or C) based on its high-order bits. EGP cannot advertise or interpret subnets or supernets.

An EGP router

- Acquires EGP neighbors
- Determines neighbor reachability
- Exchanges network reachability information with its neighbors

Each of these capabilities has an associated phase in EGP: the Neighbor Acquisition phase, the Neighbor Reachability phase, and the Network Reachability phase, respectively.

In the acquisition phase, EGP is responsible for forming neighbor relationships between routers that are peers. Routers that are peers each have an interface to a common network. One router attempts to acquire a peer router. If the peer agrees to be acquired, the two routers form a neighbor relationship. They then negotiate the mode of operation and the polling modes.

Once two routers agree to form a neighbor relationship, they must then negotiate modes. According to EGP, the routers' modes are determined as shown in Table 7-1.

Table 7-1. Router Mode Determinator

Router A	Router B	Resulting Modes
Active	Passive	Router A is active; Router B is passive.
Passive	Passive	Not allowed.
Active	Active	The router with the lower autonomous system number becomes active; the other becomes the passive router.
Both	Active	Router A is passive; Router B is active.
Both	Passive	Router A is active; Router B is passive.
Both	Both	The router with the lower autonomous system number becomes active; the other becomes the passive router.

Table 7-1 shows all possible acquisition mode combinations that are available when you configure the EGP neighbors at each end of a connection. However, it is recommended that one router be configured in the Active acquisition mode and the other in the Passive acquisition mode.

In the neighbor reachability phase, EGP is responsible for monitoring and maintaining an established EGP neighbor relationship between two routers. Its purpose is to ensure that the neighbors are operational and can provide reliable network reachability information. Two neighbors will be able to exchange network reachability information only if they are both in the UP state and know that they are both in the UP state. This is the point at which neighbor reachability is positively determined.

In the network reachability phase, EGP is responsible for determining which networks are reachable through two EGP neighbors; that is, it provides the network reachability information. This information provides a list of gateways, the networks those gateways can reach, and their associated distances.

Two neighbors determine network reachability by exchanging Poll messages and Routing Update responses as described below:

- The active neighbor sends a Poll message to a passive neighbor that it already knows to be reachable. The Poll message requests routing information from the passive neighbor.
- The Routing Update Response is the message that contains the routing information (the list of gateways on the common network, the networks they can reach, and associated distances). Both active and passive neighbors can send Routing Update messages. The active neighbor usually sends a Routing Update Response after it sends a Poll message. The passive neighbor usually sends a Routing Update Response in response to a Poll message.

EGP Implementation Notes

This section provides you with some important guidelines to follow when you configure EGP. If you do not follow these guidelines, EGP will become disabled on the interfaces involved.

- Autonomous system numbers must be between 1 and 65535.
- Two autonomous systems connected by an EGP link must have different autonomous system numbers.
- The remote IP address cannot be the same as any of the local IP interface addresses.
- The remote IP address must be on the same subnet as one of the local IP interfaces.
- EGP does not have any loop avoidance techniques -- avoid loop topologies; otherwise, you will have to configure EGP route filters to counter the redundancies.
- An EGP configuration between two ASs, each using a subnetted interface to a class A network, results in a routing black hole. RIP Version 1 aggregates the single subnet into the natural network, but the gateway router does not have complete subnet information or a natural network route to match the one being advertised by RIP Version 1. This black hole is not an EGP or RIP defect: rather, it is caused by the way RIP aggregates subnets into natural networks. EGP will operate over a subnetted interface between two Bay Networks routers if a static route is implemented. The router accepting the subnet from the remote network must augment the single subnet information with a static route for the entire remote network.

Configuring EGP on the Router

The following sections describe settings that affect the way EGP runs on the router.

Enabling and Disabling EGP

When you start EGP on the router, EGP is automatically enabled for operation.

You can use Site Manager to disable and re-enable EGP.

Site Manager: Enable parameter: page A-18

Supplying a Local AS Number

Each autonomous system has a NIC-assigned decimal number ID.

You must supply the AS ID for the local autonomous system (the AS to which this router belongs). There is no default for this parameter.

You can use Site Manager to supply the local AS number.

Site Manager: Local Autonomous System ID parameter: page A-18

Configuring a Neighbor

The following sections describe settings that define an EGP neighbor relationship.

Specifying the Neighbor's Address

You define the neighbor-to-neighbor relationship by specifying the IP address of the router that is to be the remote neighbor.

You can use Site Manager to supply the address of the remote neighbor.

Site Manager: Remote Autonomous System IP Address parameter: page A-19

Specifying the Gateway Mode

You can configure the EGP router to operate in one of two gateway modes for any given IP interface:

- Noncore

  When the router is configured as a noncore gateway, the AS to which it belongs acts as a stub AS. It advertises and forwards only traffic that originated or is destined for a network within its AS.

- Core

  When the router is configured as a core gateway, the AS to which it belongs acts as a transit AS. In the core mode, it can advertise and forward traffic to networks reachable interior or exterior to its local AS.

The default gateway mode is core mode. If the EGP router is reconfigured to run in noncore mode, the Site Manager *automatically* configures EGP export route filters on that IP interface. This is done to suppress OSPF external routes to EGP and the advertisement of any networks learned by EGP.

You can use Site Manager to specify the gateway mode of this EGP neighbor.

Site Manager: Gateway Mode parameter: page A-19

If you choose Non Core, the AS to which this EGP neighbor belongs will act as a stub AS. That is, it will only advertise networks that reside within the AS.

Enabling and Disabling the Neighbor Relationship

When you establish a neighbor-to-neighbor relationship on an interface, the relationship is automatically enabled.

You can use Site Manager to temporarily disable this neighbor relationship rather than delete it.

Site Manager: Enable parameter: page A-20

Selecting the Acquisition Mode

In an EGP neighbor relationship, one router is the active neighbor and the other router is the passive neighbor. The router in the active mode is the initiator.

By default, EGP assumes that the remote router is the passive neighbor.

You can use Site Manager to identify the remote router as the active neighbor.

Site Manager: Acquisition Mode parameter: page A-20

Selecting the Poll Mode

The type of neighbor reachability algorithm executed by the local EGP neighbor is called the poll mode. There are two poll modes: active and passive.In the active mode, a router sends Hello and Poll messages to request reachability status from its neighbor. In the passive mode, a router responds to Hello and Poll messages with I-H-U and update messages.

By default, EGP is configured to execute both the active and passive poll mode. You can use Site Manager to execute the active or passive mode only.

Site Manager: Poll Mode parameter: page A-20

Setting Neighbor Timers

EGP uses two configurable timers: the hello timer and the poll timer.

The hello timer determines the number of seconds between the local EGP neighbor's EGP Hello message retransmissions. This variable represents the RFC 904 t1 timer. By default, EGP sends a hello message every 60 seconds.

You can use Site Manger to supply a value between 30 and 120 seconds.

The poll timer determines the interval between the local EGP neighbor's EGP Poll message retransmissions. This variable represents the RFC 904 T2 timer. By default, EGP retransmits a poll message every 180 seconds. You can use Site Manager to specify a value between 120 and 480 seconds.

Site Manager: Hello Timer parameter: page A-21
Site Manager: Poll Timer parameter: page A-21

Chapter 8
Customizing NetBIOS over IP

This chapter consists of the following sections:

- NetBIOS Concepts and Terminology -- page 8-2
  - NetBIOS in an IP Environment -- page 8-3
  - Forwarding Name Queries over an Unnumbered Interface -- page 8-6
- Enabling and Disabling NetBIOS on a Router -- page 8-6
- Specifying a TTL Value for a Rebroadcast Packet -- page 8-6
- Enabling the Insertion of Record Route Option -- page 8-7
- Configuring a NetBIOS Cache -- page 8-7
  - Enabling Name Caching on the Router -- page 8-8
  - Creating a MIB Instance for a Cached Name -- page 8-8
  - Specifying the Size of the Name Cache -- page 8-8
  - Aging a Cache Entry -- page 8-9
  - Customizing a Cache Search -- page 8-10
- Configuring NetBIOS on an IP Interface -- page 8-11
  - Enabling and Disabling NetBIOS on an Interface -- page 8-11
  - Enabling Name Caching on the Interface -- page 8-11
  - Disabling Inbound and Outbound Broadcasts -- page 8-11
  - Supplying a Rebroadcast Address -- page 8-12
- Configuring a Static NetBIOS Name and Address -- page 8-12
  - Enabling and Disabling Static Name Caching -- page 8-12
  - Specifying the NetBIOS Station Name -- page 8-12
  - Specifying the NetBIOS Station Address -- page 8-13
- Adding a Traffic Filter to a NetBIOS Interface -- page 8-14

NetBIOS Concepts and Terminology

The Network Basic Input-Output System (NetBIOS) is a session layer communications service used by client and server applications in IBM Token Ring and PC LAN networks.

NetBIOS provides applications with a programming interface for sharing services and information across a variety of lower-layer network protocols, including IP. Figure 8-1 shows the position of NetBIOS and IP in a simple network architecture.

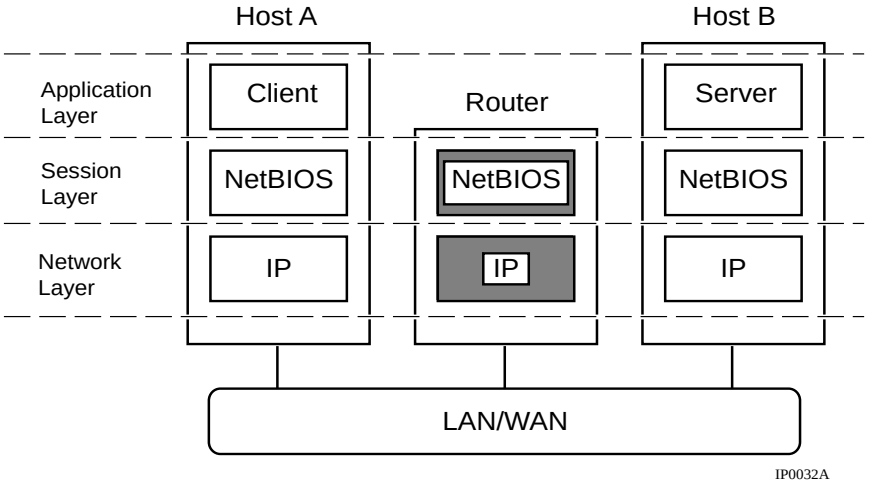

Figure 8-1. NetBIOS over IP

There are three categories of NetBIOS services: the name service, the session service, and the datagram service.

The NetBIOS *name service* allows an application to

- Verify that its own NetBIOS name is unique. The application issues an Add Name Query to NetBIOS. NetBIOS broadcasts the Add Name Query, containing the name. NetBIOS applications that receive the query return an Add Name Response or a Name in Conflict Response. If no response to the query is received after (typically) six broadcasts, the name is considered to be unique.
- Delete a NetBIOS name that the application no longer requires.

- Use a server's NetBIOS name to determine the server's network address. The application issues a Name Query Request to NetBIOS, containing the target server's NetBIOS name. NetBIOS broadcasts the Name Query Request. The server that recognizes the name returns a Name Query Response containing its network address.

The NetBIOS *session service* allows an application to conduct a reliable, sequenced exchange of messages with another application. The messages can be up to 131,071 bytes long.

The NetBIOS *datagram service* allows an application to exchange datagrams with a specific application or to broadcast datagrams to a group and receive datagrams from the group. Datagrams allow applications to communicate without establishing a session. When a NetBIOS application wants to send information that does not require acknowledgment from the destination application, the application can transmit a NetBIOS datagram.

NetBIOS in an IP Environment

The NetBIOS name service and datagram service rely on the capability of the underlying network to broadcast Name Query Requests to all NetBIOS applications. In a NetBIOS over IP environment, it is the responsibility of the IP router to ensure that the broadcast queries reach all appropriate network segments. To do this, the router

1. Analyzes each NetBIOS packet received on any NetBIOS interface to determine whether the packet is a broadcast packet
2. Rebroadcasts each broadcast packet out all appropriate interfaces except the one on which it was received (readdressing the packet if required)

If alternate paths exist between different network segments, broadcasting loops can occur. To prevent such loops, the router

1. Stamps the data portion of the IP packet with the IP address of the router from which the packet was rebroadcast
2. Parses the IP addresses included in the data portion of the IP packet to determine if the packet has already been rebroadcast by that router

In Figure 8-2, for example, Client C on the network connected to Router B wishes to communicate with Server S, which is located on the network connected to Router A.

1. The client issues a Name Query Request to NetBIOS on the host, specifying the server application by its NetBIOS name. The IP service on the host broadcasts the Name Query Request.
2. Router B receives the Name Query Request, determines that it is a broadcast message, and rebroadcasts it out each of its NetBIOS interfaces (except for the one on which it arrived).
3. Router A receives the broadcast Request and rebroadcasts to its local network.
4. The server on Router A receives the IP broadcast Request and recognizes its own name.

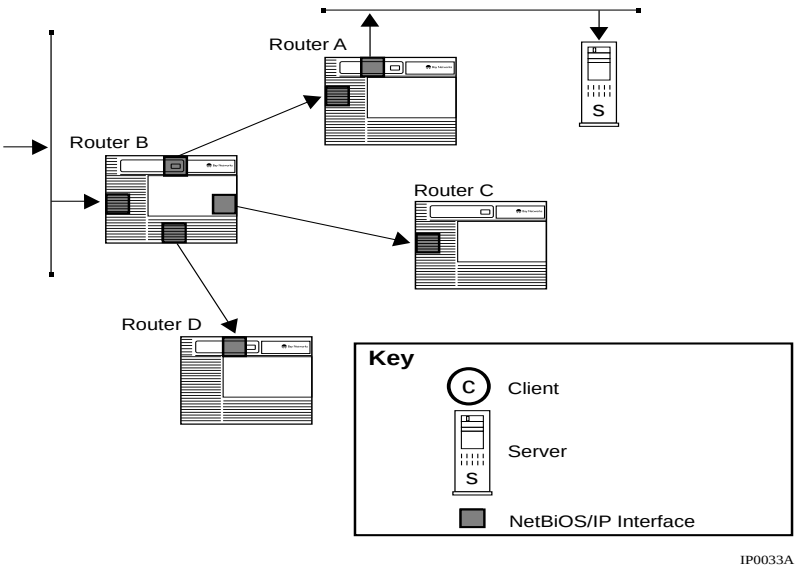

Figure 8-2. Broadcasting a Name Query Request

The server responds to the Name Query Request by issuing a Positive Name Query Response, containing the IP address of the server, to NetBIOS on the host. The following steps occur (Figure 8-3):

1. NetBIOS sends the Response to Router A as a unicast message.
2. Router A and Router B forward the unicast Response to the awaiting client**.**

Now that the client has obtained the server's IP address from the Name Query Response, client and server can communicate by exchanging IP messages.

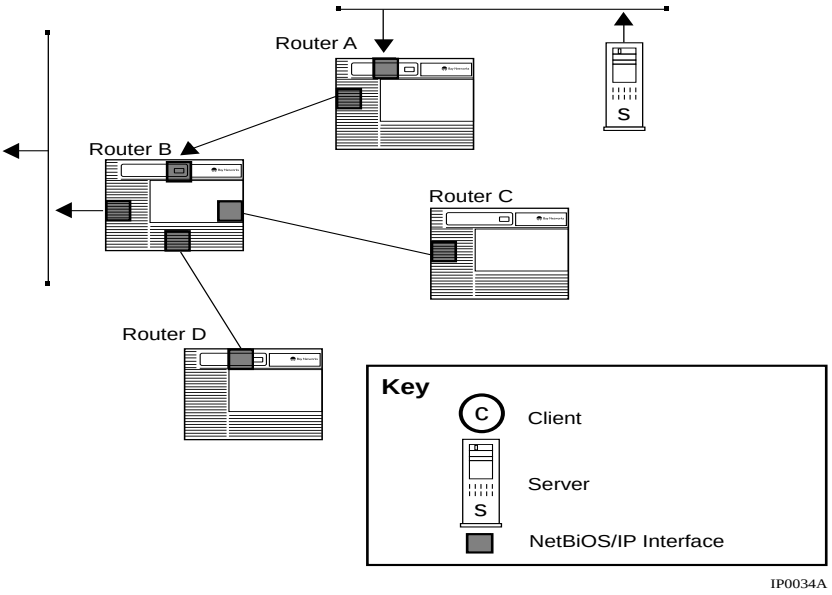

Figure 8-3. Returning a Unicast Name Query Response

Forwarding Name Queries over an Unnumbered Interface

NetBIOS cannot be configured directly on an unnumbered interface. Because of this restriction, Name Query Requests cannot be broadcast over an unnumbered interface.

To forward Name Query Requests over unnumbered interfaces, the network administrator configures a static NetBIOS name entry to the proper NetBIOS name server (see “Enabling and Disabling Static Name Caching” on page 8-12). In this way, Name Query Requests will traverse the unnumbered interfaces as unicast IP packets.

Enabling and Disabling NetBIOS on a Router

When you start NetBIOS on the router, NetBIOS is automatically enabled.

You can use Site Manager to disable and re-enable NetBIOS as required.

Site Manager: Enable/Disable parameter: page A-66

Specifying a TTL Value for a Rebroadcast Packet

The TTL value restricts the number of routers a rebroadcast packet can traverse. To prevent NetBIOS broadcast packets from traversing the network indefinitely, set the parameter to a minimal value.

By default, NetBIOS sets the TTL value in each packet to 5. You can use Site Manager to set a TTL value from 1 to 255.

Site Manager: Rebroadcast Packet TTL parameter: page A-69

Enabling the Insertion of Record Route Option

The Insertion of Record Route option prevents rebroadcast packets from looping forever by allowing the NetBIOS entity in the router to determine whether it has received this packet before on this interface. If so, the router drops it.

By default, the NetBIOS Insertion of Record Route option in rebroadcast packets is disabled. You can use Site Manager to enable the option.

Site Manager: Rebroadcast Record Route parameter: page A-69

If all IP entities support this option, enable it on the router.

Configuring a NetBIOS Cache

NetBIOS is a broadcast-intensive protocol. Much of the broadcast overhead is related to maintaining unique names across the network and providing end users with access to NetBIOS applications. The amount of overhead grows with the number of NetBIOS resources (applications, servers, and clients) on the network.

To keep broadcast traffic to a minimum, each router that runs NetBIOS over IP builds and maintains a cache of NetBIOS name/IP address pairs, using information contained in the Name Query Responses it receives and forwards.

In Figure 8-3, for example:

1. Router A receives a Name Query Response from the server. The router gleans from the Name Query Response the name and IP address of the server.
2. The router stores the name and IP address of the server in its cache.
3. The router forwards the Name Query Response.

Routers that support NetBIOS must analyze each Name Query Request received on a NetBIOS interface to determine whether the name of the requested resource (typically, a server) is in the cache. If so, the router replaces the broadcast address in the Request with the unicast IP address of the server. The router then forwards the Name Query Request to the server.

Enabling Name Caching on the Router

NetBIOS name caching gives the router the ability to cache the name associated with each NetBIOS server that is active on the network.

By default, NetBIOS name caching is disabled. You can use Site Manager to enable name caching.

15-character NetBIOS name caching gives the router the ability to treat a NetBIOS name as either a 15- or a 16-character entity.

By default, NetBIOS treats a name as a 16-character entity. You can use Site Manager to enable the feature if you want NetBIOS to treat a name as a 15-character entity.

Site Manager: NetBIOS Name Caching parameter: page A-66
Site Manager: 15-Character NetBIOS Name Caching parameter: page A-66

Creating a MIB Instance for a Cached Name

By default, NetBIOS creates a MIB instance for each name entry stored in the name cache.

You can use Site Manager to disable the feature if you want to release the system memory and processing resources otherwise dedicated to maintaining cached names in the MIB.

Site Manager: Create MIB Inst for Cached Name parameter: page A-67

Specifying the Size of the Name Cache

By default, NetBIOS allocates space for 100 entries in the name cache. You can accept the default or use Site Manager to specify a value of 1 to 2147483647 entries.

Site Manager: Max Name Cache Entries parameter: page A-67

You can adjust this value in direct proportion to the total number of server names expected to be active during intervals of peak traffic load or performance demand on the router. A value of 100 is suitable for networks that include up to 100 NetBIOS names to cache.

Aging a Cache Entry

The router ages cache entries to ensure that cached routes remain consistent with the current network topology. If the cache table lookup mechanism does not access a cache entry within the period you set in the appropriate Cache Aging Time parameter, the router deletes the entry from the table.

If the router receives a broadcast Name Query Request from a client and finds the name and associated IP address of the requested server in its cache, the router replaces the broadcast address on the Name Query Request with the unicast IP address. The router also assigns the entry a short time to live. If the entry is valid, the router will receive a Positive Name Query Response (which will validate the entry) from the server within the specified time to live. If the entry is invalid, the Name Query Request will not reach the server. In this case, the entry quickly ages out.

By default, inactive NetBIOS names expire from the NetBIOS name cache after 300 seconds.

You can use Site Manager to specify any time value that can rapidly age infrequently referenced names out of the NetBIOS name cache.

Site Manager: Name Cache Age parameter: page A-68

Choose an aging value that allows infrequently referenced or obsolete server names to expire from the name cache. The smaller the value, the less efficient broadcast reduction is, but the more quickly the network recovers topology changes.

Customizing a Cache Search

The mechanism that NetBIOS uses to search for a name in the cache is based on a fast string hash/search mechanism developed for AppleTalk Zone Name processing. This mechanism uses a hash table that NetBIOS builds and maintains on the router.

Increasing the number of entries in the hash table

- Decreases the number of names the router must compare before finding a specific cached name
- Decreases the amount of time it takes the router to find a particular cached name
- Increases memory usage

Note that increasing the number of entries in the hash table does not increase the number of names the router can cache. This is determined by user configuration and by available memory.

By default, NetBIOS allows 253 names in the hash table.

You can use Site Manager to specify the number of entries you want to allow in the cache lookup tables.

Site Manager: Hash Entry Count parameter: page A-68

For networks that actively use up to 2500 NetBIOS server names, use the default value (253). To determine a hash entry count for larger networks:

- Divide the total number of unique NetBIOS server names active in the network by 10.
- Adjust the quotient to the nearest (higher or lower) prime number. (A prime number can only be divided by itself or by 1 and still yield a whole-number quotient.)
- Replace the default value with the new, calculated number.

Increasing the number of hash table entries does not increase the number of names that a router can cache. With larger networks, increasing the size of the hash tables may, however, reduce internal cache lookup time, thereby improving overall performance.

Configuring NetBIOS on an IP Interface

NetBIOS software on the router communicates with NetBIOS clients and servers through IP interfaces that have been configured with NetBIOS.

Enabling and Disabling NetBIOS on an Interface

When you configure NetBIOS on an interface, NetBIOS is automatically enabled. You can use Site Manager to disable and re-enable NetBIOS on the interface.

Site Manager: Enable/Disable parameter: page A-69

Enabling Name Caching on the Interface

By default, NetBIOS name caching is enabled on the interface.

You can use Site Manager to disable the ability of this interface to cache the name of each NetBIOS server active in the network.

Site Manager: NetBIOS Name Caching parameter: page A-70

Select Disable if you want to release system memory and processing resources otherwise dedicated to server name caching.

Disabling Inbound and Outbound Broadcasts

By default, NetBIOS can receive inbound broadcasts on the interface and send outbound broadcasts. You can use Site Manager to disable this feature on the interface.

Site Manager: Enable NetBIOS Inbound Broadcasts parameter: page A-70
Site Manager: Enable NetBIOS Outbound Broadcasts parameter: page A-70

Supplying a Rebroadcast Address

By default, NetBIOS uses the IP broadcast address configured for this interface when rebroadcasting NetBIOS packets out this interface. You can use Site Manager to supply a rebroadcast address that overrides this broadcast address.

Site Manager: Rebroadcast Address parameter: page A-71

Configuring a Static NetBIOS Name and Address

You can add static NetBIOS names to the router. These entries are independent of the name entries learned dynamically in the name cache.

Enabling and Disabling Static Name Caching

By default, NetBIOS caches the names you have added statically. You can use Site Manager to disable this feature.

Site Manager: Enable parameter: page A-71

Specifying the NetBIOS Station Name

You can use Site Manager to specify the 1- to 16-character name of a NetBIOS station.

Site Manager: NetBIOS Station Name parameter: page A-72

The name must not exceed 16 characters. The system pads names shorter than 16 characters with ASCII space characters.

Specifying the NetBIOS Station Address

You can use Site Manager to specify an IP address to associate with the statically configured name of the NetBIOS station.

Site Manager: IP Address parameter: page A-72

Supplying a Scope Identifier

When you configure a static name, you must specify its NetBIOS scope -- that is, the area of the network across which the name is known. Each NetBIOS scope has a Scope Identifier, a string of characters that meet the requirements of the Domain Name System. (All NetBIOS names are represented in a manner consistent with the definition for “compressed name messages” outlined in the Domain Name Service Specification, RFC 833.)

Site Manager: NetBIOS Scope ID parameter: page A-71

Adding a Traffic Filter to a NetBIOS Interface

If name caching is enabled, a router that receives a Name Query Response (originating from a server and addressed to a client) must be able to deliver the message to the NetBIOS entity on the router (rather than simply forward it out another interface toward its destination).

To enable the router to recognize a unicast IP packet that contains a Name Query Response and pass it to NetBIOS through UDP port 137, you must configure a traffic filter on each NetBIOS interface that receives unicast Name Query Responses.

1. **Beginning at the Configuration Manager window, select Circuits** > **Edit.**

   The Circuit List window opens.

2. **Click on Edit.**

   The Circuit Definition window opens.

3. **Select Protocols** > **Edit IP** > **Traffic Filters.**

   The IP Filters window opens.

4. **Click on Template.**

   The Filter Template Management window opens.

5. **Click on Create.**

   The Create IP Template window opens.

6. **Select Criteria** > **Add** > **UDP Frame** > **Destination Port.**

   The Edit Range screen opens.

7. **Enter 137 for the minimum value and the maximum value.**

8. **Click on OK.**

   The Create IP Template window opens.

9. **Select Action** > **Add** > **Forward to Next Hop.**

   The Next Hop window opens.

10. **Enter the IP address of this interface (the interface on which you are configuring the traffic filter). Click on OK.**

Appendix A
Site Manager Parameters

BGP Parameters

BGP Configuration Parameters

Parameter:	**Identifier**
Path:	Select BGP on the Select protocols window.
Default:	None
Options:	An IP address of an IP interface on this router
Function:	Identifies the BGP router. There is no default for this parameter. You must use an IP address of one of the router's IP interfaces.
Instructions:	Either accept the current BGP identifier or enter a new IP address. The BGP identifier must be one of the router's IP interfaces. If both BGP and OSPF are running on the router, then the OSPF router ID must be equivalent to one of the configured IP interfaces.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.1.4

Parameter:	**Local AS**
Path:	Select BGP on the Select protocols window
Default:	None
Options:	1 to 65535
Function:	Identifies the autonomous system to which this BGP router belongs.
Instructions:	Enter a value from 1 to 65535.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.1.5

BGP Global Parameters

Parameter: BGP Enable

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: Enable
Options: Enable | Disable
Function: Globally enables or disables BGP on all router interfaces.
Instructions: Set to Disable if you want to disable BGP for the entire router. Set to Enable if you previously disabled BGP and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.2

Parameter: BGP Identifier

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: None
Options: An IP address of an IP interface on this router
Function: Identifies the BGP router. There is no default for this parameter. You must use an IP address of one of the router's IP interfaces.
Instructions: Either accept the current BGP identifier or enter a new IP address. The BGP identifier must be one of the router's IP interfaces. If both BGP and OSPF are running on the router, then the OSPF router ID must be equivalent to one of the configured IP interfaces.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.4

Parameter: BGP Local AS

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: None
Options: 1 to 65535
Function: Identifies the autonomous system to which this BGP router belongs.
Instructions: Either accept the current BGP Local AS value or enter a new value for this parameter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.5

Parameter: BGP Intra-AS

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: Enable
Options: Enable | Disable
Function: Specifies whether BGP will perform intra-AS IBGP routing.
Instructions: Transit ASs should use intra-AS routing. Stub or multihomed ASs usually do not use IBGP intra-AS routing.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.8

Parameter: BGP From Protocols

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: BGP
Options: BGP | ALL
Function: Controls (if intra-AS routing is enabled) the types of routes that BGP advertises in any IBGP sessions.
Instructions: Select BGP to propagate only advertised routes learned from external BGP peers. Select ALL to propagate routes learned from all route sources (excluding IBGP and OSPF interarea and intra-area routes, which are never advertised with IBGP).
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.9

Parameter: BGP Interval Timer

Path: Configuration Manager > Protocols > IP > BGP > Global
Default: 5 seconds
Options: 1 to 2147483647
Function: Specifies the minimum time interval, in seconds, between injections of external BGP routes into the IP routing table.
Instructions: Accept the default or enter a nonzero value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.10

Parameter: BGP Collision Detect

Path: Configuration Manager > Protocols > IP > BGP > Global

Default: Enable

Options: Enable │ Disable

Function: Specifies whether redundant BGP connections to the same router will be detected and disallowed. If you want only one BGP connection to the same router to be maintained, use the default. If you want to allow redundant connections, enter Disable.

Instructions: Collision detection is based on router ID. If two BGP peers have multiple physical connections and want to establish a BGP session across each physical connection, you must disable this parameter. The advantage of a configuration with multiple physical connections is redundancy. The disadvantage is that such a configuration results in multiple copies of each route.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.16

Parameter: Multi-hop EBGP Connection

Path: Configuration Manager > Protocols > IP > BGP > Global

Default: Disable

Options: Enable │ Disable

Function: Specifies whether BGP allows multihop connections to an external BGP peer.

Instructions: By default, BGP enforces the rule that requires an external BGP peer to be located on a directly attached network. Use this parameter to override the restriction. Enabling multihop BGP connections is dangerous because it can cause BGP speakers to establish a BGP connection that traverses a third-party AS, which may violate policy considerations and may also introduce forwarding loops.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.6

Parameter: **Dynamic Policy Change**

Path: Configuration Manager > Protocols > IP > BGP > Global

Default: Enable

Options: Enable │ Disable

Function: Specifies whether or not BGP dynamically re-evaluates all routes affected by a policy when you modify the policy, .

Instructions: Select disable if you want BGP to restart all connections when you modify a policy.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.19

Parameter: **Soloist Slots**

Path: Configuration Manager > Protocols > IP > BGP > Global

Default: All slots

Options: Slots 1 through 14

Function: Specifies slots on which the BGP soloist is eligible to run.

Instructions: Use the ISP Mode parameter (IP global) to configure BGP as a soloist.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.20

BGP-3 Global Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > BGP > BGP-3 > Global

Default: Enable

Options: Enable │ Disable

Function: Globally enables or disables BGP-3 on all router interfaces.

Instructions: Set to Disable if you want to disable BGP-3 for the entire router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.2.1.2

BGP-4 Global Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > BGP > BGP-4 > Global
Default: Enable
Options: Enable | Disable
Function: Globally enables or disables BGP-4 on all router interfaces.
Instructions: Set to Disable if you want to disable BGP-4 for the entire router.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.3.1.2

BGP Peer Parameters

Parameter: **Peer Address**

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: None
Options: Any IP address
Function: Specifies the IP address of the interface on the remote side of this BGP peer connection.
Instructions: Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.6

Parameter: **Peer AS**

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: None
Options: 1 to 65535
Function: Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs.
Instructions: Enter the appropriate AS number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.10

Parameter: Local Address

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: None
Options: Any IP address
Function: Specifies the IP address of the interface on the local side of this BGP peer connection.
Instructions: Enter the appropriate address.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.4

Parameter: Enable

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: Enable
Options: Enable | Disable
Function: Enables or disables a BGP peer relationship with the specified IP address.
Instructions: Set this parameter to Disable if you want to temporarily disable this peer relationship rather than delete it. Or set it to Enable if you previously disabled this peer relationship and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.2

Parameter: Min BGP Version

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: 4
Options: 3 or 4
Function: Specifies the minimum acceptable BGP version to run on this peer connection.
Instructions: Specify BGP-3 or BGP-4.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.8

Parameter: Max BGP Version

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: 4
Options: 3 or 4
Function: Specifies the maximum acceptable BGP version to run on this peer connection.
Instructions: Specify BGP-3 or BGP-4.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.9

Parameter: Peer AS

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: None
Options: 1 to 65535
Function: Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs.
Instructions: Either accept the current value or enter a new one.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.10

Parameter: External Advertisement Timer

Path: Configuration Manager > Protocols > IP > BGP > Peers
Default: 5 seconds
Options: 1 to 2147483647
Function: Specifies the minimum number of seconds allowed between BGP updates for this peer connection.
Instructions: Either accept the current value or enter a value greater than zero seconds.
MIB Object ID: The external advertisement interval controls how often the IP routing table is examined for changes. BGP update messages for routes that originate external to this AS will be issued no faster than the number of seconds you specify with this parameter.

Parameter: Connect Retry Timer

Path:	Configuration Manager > Protocols > IP > BGP > Peers
Default:	120 seconds
Options:	0 to 2147483647
Function:	Specifies the maximum number of seconds allowed between TCP connection attempts for this peer connection.
Instructions:	Either accept the current value or set this parameter to some other value. A value of 0 indicates that no active attempt to establish a BGP connection to the peer is to be done. Incoming calls from the peers will be accepted.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.12

Parameter: Holdtime

Path:	Configuration Manager > Protocols > IP > BGP > Peers
Default:	90 seconds
Options:	0 or any decimal number greater than 2
Function:	Specifies the holdtime that will be inserted into an Open message. Upon receipt of the peer's Open message, the lesser of the two holdtimes will be used (this must be at least 3 seconds). There are two exceptions. If one peer sends a zero holdtime, then the non-zero holdtime is used. If both peers send zero holdtimes, then no holdtime is used. The calculated holdtime is the amount of time either peer will wait for a Keepalive or Update message before declaring the connection down.
Instructions:	Either accept the current Holdtime Timer value or set the parameter to 0 or some value greater than 2 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.13

Parameter: Keepalive Timer

Path: Configuration Manager > Protocols > IP > BGP > Peers

Default: 30 seconds

Options: Any decimal number

Function: Specifies how often Keepalive messages will be sent across this peer connection.

Instructions: If a holdtime of 0 is negotiated, no periodic Keepalive messages are sent. Otherwise, the Keepalive timer is set to the smaller of this configured value and one-third of the holdtime.

MIB Object ID: Either accept the current Keepalive value or set this parameter to some value greater than 0.

Parameter: Min AS Origination Interval

Path: Configuration Manager > Protocols > IP > BGP > Peers

Default: 15 seconds

Options: A value greater than 0

Function: Determines the minimum amount of time that must elapse between successive advertisements of Update messages that report changes within the advertising BGP speaker's own autonomous system.

Instructions: Enter a value greater than 0 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.30

Parameter: Local AS to Advertise to Peer

Path: Configuration Manager > Protocols > IP > BGP > Peers

Default: Null

Options: 1 to 65535

Function: Specifies the AS number that is sent in an Open message to this peer.

Instructions: Enter an AS number. To specify the AS number you set with the BGP Local AS parameter, use the default, null.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.31

Parameter: Peer Max Update Size

Path: Configuration Manager > Protocols > IP > BGP > Peers

Default: 800 bytes

Options: 64 to 4096 bytes

Function: Specifies the maximum size (in bytes) of Update messages that are sent to this peer.

Instructions: Use the default or specify a size. Note that, if the size of the Update message that is used to advertise a single route is greater than the configured message size, the actual message size can exceed the configured value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.32

Parameter: Peer Route Echo Switch

Path: Configuration Manager > Protocols > IP > BGP > Peers

Default: Enable

Options: Enable | Disable

Function: Controls the way the router echoes a BGP route that is selected for forwarding. (Echoing in this case means advertising the route back to the peer from which it was received.) If this parameter is enabled, the router advertises the route back as reachable and includes the local AS. If this parameter is disabled, the router echoes the route as UNREACHABLE/withdrawn.

Instructions: If the peer router saves routes that contain its own AS number and is running short of memory, send an UNREACHABLE echo.

MIB Object ID: A BGP speaker that participates in inter-AS multicast routing must advertise a route it receives from one of its external peers. If the router stores the route in its routing table, it must also advertise it back to the peer from which the route was received. For a BGP speaker that does participate in inter-AS multicast routing, such echoing is optional.

BGP AS Weight and Weight Class Parameters

Parameter: **AS**

Path: Configuration Manager > Protocols > IP > BGP > Weights
Default: Null
Options: 1 to 65535
Function: Identifies the autonomous system to which you want to assign a weight.
Instructions: Enter the appropriate AS number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.4

Parameter: **Weight Value 1**

Path: Configuration Manager > Protocols > IP > BGP > Weights
Default: 8
Options: 1 to 15, plus the infinity value of 16
Function: Specifies the Class 1 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.
Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: **Weight Value 2**

Path: Configuration Manager > Protocols > IP > BGP > Weights
Default: 8
Options: 1 to 15, plus the infinity value of 16
Function: Specifies the Class 2 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.
Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 3

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 3 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 4

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 4 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 5

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 5 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 6

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 6 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 7

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 7 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 8

Path: Configuration Manager > Protocols > IP > BGP > Weights

Default: 8

Options: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 8 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > BGP > Weights
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables a weight assignment for a particular AS.
Instructions:	Set to Disable to disable the weight assignment for this AS; set to Enable if you previously disabled this weight assignment and now want to re-enable it.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.2

BGP Event Message Parameters

Parameter: Local IP Address

Path: Configuration Manager > Protocols > IP > BGP > Debug
Default: Null
Options: An IP address
Function: Specifies a BGP peer's local address.
Instructions: Enter 0.0.0.0. to obtain event messages about all connections to a peer with the specified local address.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.2

Parameter: Remote Address

Path: Configuration Manager > Protocols > IP > BGP > Debug
Default: Null
Options: An IP address
Function: Specifies a BGP peer's remote address.
Instructions: Enter 0.0.0.0 to obtain event messages about all connections to peers using the specified remote address.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.3

Parameter: Message Level

Path: Configuration Manager > Protocols > IP > BGP > Debug
Default: ALL
Options: ALL | DEBUG | INFO | WARNING | FAULT | TRACE
Function: Specifies the severity level of event messages required.
Instructions: Select the default to obtain event messages of all levels.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.4

Parameter:	**Message Trace Switch**
Path:	Configuration Manager > Protocols > IP > BGP > Debug
Default:	DISABLE
Options:	DISABLE │ OPEN │ UPDATE │ NOTIFICATION │ KEEPALIVE
Function:	Specifies whether or not BGP messages on the specified connection are logged and, if so, which messages are logged.
Instructions:	Use the default or select a BGP message type.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.5

EGP Parameters

EGP Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > EGP
Default: Enable
Options: Enable │ Disable
Function: This parameter allows you to globally enable or disable EGP on all router interfaces.
Instructions: Set to Disable if you want to disable EGP for the entire router. Set to Enable if you previously disabled EGP and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.1.2

Parameter: Local Autonomous System ID

Path: Configuration Manager > Protocols > IP > EGP
Default: None
Options: 1 to 65535
Function: Identifies the local autonomous system (the AS to which this router belongs) by the NIC-assigned decimal number. There is no default for this parameter.
Instructions: Either accept the current value for this parameter or enter a new value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.1.7

EGP Neighbor Parameters

Parameter:	**Remote Autonomous System IP Address**
Path:	Configuration Manager > Protocols > IP > EGP > Neighbors
Default:	Null
Options:	Any IP address
Function:	Specifies the IP address of the remote router that will form an EGP neighbor relationship with this router.
Instructions:	Enter the IP address in dotted decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.4.3.1.4

Parameter:	**Gateway Mode**
Path:	Configuration Manager > Protocols > IP > EGP > Neighbors
Default:	Core
Options:	Core | Non Core
Function:	Specifies the gateway mode for this EGP neighbor. If you choose Core, the default, the local AS to which this EGP neighbor belongs will act as a transit AS. That is, it will advertise networks that reside within the AS as well as within external networks.
Instructions:	If you choose Non Core, the AS to which this EGP neighbor belongs will act as a stub AS. That is, it will only advertise networks that reside within the AS. Set this parameter to either Core or Non Core, depending on how you want this EGP neighbor to function.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.4.3.1.5

Parameter: Enable

Path: Configuration Manager > Protocols > IP > EGP > Neighbors
Default: Enable
Options: Enable | Disable
Function: Enables or disables an EGP neighbor relationship with the specified IP address.
Instructions: Set this parameter to Disable if you want to temporarily disable this neighbor relationship rather than delete it. Or set it to Enable if you previously disabled this neighbor relationship, and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.3.1.2

Parameter: Acquisition Mode

Path: Configuration Manager > Protocols > IP > EGP > Neighbors
Default: Passive
Options: Passive | Active
Function: Specifies which of the two neighbors initiates EGP connections. The router in the active mode is the initiator.
Instructions: Set this parameter to Active if you want the local EGP neighbor to be the initiator of EGP connections. Otherwise, accept the default value, Passive.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.3.1.7

Parameter: Poll Mode

Path: Configuration Manager > Protocols > IP > EGP > Neighbors
Default: Both
Options: Active | Passive | Both
Function: Specifies the type of neighbor reachability algorithm this local EGP neighbor executes. In the active mode, a router sends Hello and Poll messages to request reachability status from its neighbor. In the passive mode, a router responds to Hello and Poll messages with I-H-U and update messages.
Instructions: Accept the default value, Both, or set to either Active or Passive (depending on the neighbor reachability algorithm you want this router to execute).
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.3.1.8

Parameter: Hello Timer

Path:	Configuration Manager > Protocols > IP > EGP > Neighbors
Default:	60 seconds
Options:	30 to 120 seconds
Function:	Specifies the number of seconds between the local EGP neighbor's EGP Hello message retransmissions. This variable represents the RFC 904 t1 timer.
Instructions:	Accept the default value of 60 seconds for this parameter or set it to some value between 30 and 120 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.4.3.1.9

Parameter: Poll Timer

Path:	Configuration Manager > Protocols > IP > EGP > Neighbors
Default:	180 seconds
Options:	120 to 480 seconds
Function:	Specifies the time period, in seconds, between the local EGP neighbor's EGP Poll message retransmissions. This variable represents the RFC 904 T2 timer.
Instructions:	Either accept the default value of 180 seconds for this parameter or set it to some value between 120 and 480 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.4.3.1.10

IP Parameters

IP Configuration Parameters

Parameter: **IP Address**

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: 0.0.0.0 or any valid IP address

Function: Assigns a 32-bit IP address to the interface.

Instructions: Enter the IP address of the interface in dotted decimal notation. Enter 0.0.0.0 to configure an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.4

Parameter: **Subnet Mask**

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: The Configuration Manager automatically calculates an appropriate subnet mask, depending on the class of the network to which the interface connects. However, you can change the subnet mask with this parameter.

Function: Specifies the network and subnetwork portion of the 32-bit IP address.

Instructions: Either accept the assigned subnet mask or enter another subnet mask in dotted decimal notation. Enter 0.0.0.0 if you are configuring an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Transmit Bcast Addr

Path: Select IP from the Select Protocols window and click on OK.

Default: 0.0.0.0

Options: 0.0.0.0 or any valid IP broadcast address

Function: Specifies the broadcast address that this IP subnet uses to broadcast packets.Accepting 0.0.0.0 for this parameter specifies that the IP router will use a broadcast address with a host portion of all 1s. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you have IP address 123.1.1.1 and a subnet mask of 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. To set the explicit broadcast address of all 1s, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted decimal notation.If you set the IP Address parameter to 0.0.0.0 (to configure an unnumbered interface), Site Manager automatically sets this parameter to 255.255.255.255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8

Parameter: UnNumbered Assoc Address

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: Any valid IP address

Function: Specifies an address that IP uses when sourcing a packet. RIP uses this address to make decisions about advertising subnets over the unnumbered interface. RIP advertises subnets over the unnumbered interface if the subnets have the same mask as the associated address.

Instructions: Specify the address of any numbered interface on the router. If you are running RIP over the unnumbered interface and if you are using a subnet address as the associated address, the local and remote associated address should have the same network number. If you configure a local and remote associated address using different network numbers, you must use RIP2 mode.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.110

IP Interface Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Interfaces
Default: Enable
Options: Enable | Disable
Function: Enables or disables IP routing on this interface.
Instructions: Set to Disable to disable IP routing over this circuit.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.2

Parameter: Subnet Mask

Path: Configuration Manager > Protocols > IP > Interfaces
Default: You specified the subnet mask when you added IP to the circuit.
Options: Depend on the class of the network to which the interface connects
Function: Specifies the network and subnetwork portion of the 32-bit IP address.
Instructions: Enter the subnet mask in dotted decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Broadcast Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: You specified the broadcast address when you added IP to the circuit.

Options: 0.0.0.0 or any IP address

Function: Specifies the broadcast address that the IP router uses to broadcast packets. Accepting 0.0.0.0 for the broadcast address specifies that the IP router will use a broadcast address with a host portion of all 1s. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you have set the IP address to 123.1.1.1 and the subnet mask to 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. For the explicit broadcast address of all 1s, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.9

Parameter: Interface Cost

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 1

Options: 1 to the value of the RIP diameter (maximum 127)

Function: Sets the cost of this interface. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets transmitted out other interfaces.

Instructions: Enter the interface cost value (standard RIP implementation assigns a cost of 1); however, keep in mind that increasing this value causes the upper bound set by the RIP Network Diameter parameter to be attained more rapidly.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8

Parameter: MTU Discovery

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: On | Off

Function: Specifies whether the Reply MTU option (option 11 in RFC 1063) is enabled on this interface. When the option is enabled, this interface responds to Probe MTUs (option 12 in RFC 1063). A Probe MTU requests the minimum MTU (maximum transmission unit) of all networks an IP datagram must traverse from source to destination. By enabling this interface to respond to Probe MTUs, you eliminate transit fragmentation and destination reassembly for datagrams destined for this interface and, therefore, decrease network load.

Instructions: Select On to enable the Reply MTU option on this interface; select Off to disable the option on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.10

Parameter: Addr Mask Reply

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: On | Off

Function: Specifies whether this interface generates ICMP (Internet Control Message Protocol) address-mask-reply messages in response to valid address-mask-request messages. The interface generates ICMP address-mask-reply messages in compliance with the relevant sections of RFCs 950 and 1009.

Instructions: Select On to enable ICMP address-mask-reply message generation on this interface. Select Off to disable ICMP address-mask-reply message generation on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.11

Parameter: **All Subnet Bcast**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: On │ Off

Function: Specifies whether or not the IP router floods ASB datagrams it receives out this interface. An ASB datagram has a destination address equal to the broadcast address for an entire network (all subnets). For example, if a network interface serves the subnet 128.10.2.1 with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram.

Instructions: Specify On if you want the IP router to flood ASBs out this interface; specify Off to restrict the router from flooding ASBs out this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.12

Parameter: **Address Resolution**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: ARP

Options: ARP │ X.25_DDN │ X.25_PDN │ INARP │ ARPINARP │ NONE │ X.25 BFEDDN │ PROBE │ ARPPROBE

Function: Indicates the address resolution scheme for this interface. The default option, ARP, enables ARP on this interface. The option INARP (Inverse ARP) enables the address resolution for Frame Relay interfaces. It is used to discover the IP address of the station at the remote end of the virtual circuit. The PROBE option enables HP Probe for Ethernet interfaces.

Instructions: Depending on your network requirements, select INARP only when all Frame Relay stations support Inverse ARP. Select ARPINARP for your Frame Relay interfaces. ARPINARP enables both ARP and Inverse ARP. Select X.25_DDN for your X.25 DDN interfaces. Select X.25_PDN for your X.25 PDN interfaces. Select PROBE to enable HP Probe on the interface. Select ARPPROBE to enable both ARP and HP Probe.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.13

Parameter: Proxy

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: On | Off

Function: Specifies whether this interface uses Proxy ARP to respond to ARPs for a remote network.

Instructions: Select On to enable Proxy ARP on this interface. In order to enable Proxy ARP, you must have set the ARP parameter to Enable for this interface. When you enable Proxy ARP, the IP router assumes responsibility for IP datagrams destined for the remote network. To enable Proxy ARP for subnets reachable via a default route, also set the Enable Default Route for Subnets parameter to Enable. To enable Proxy ARP for remote destinations on other networks, set the Nonlocal ARP Destination parameter to Accept. Select Off to disable Proxy ARP on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.14

Parameter: Host Cache

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: Off | 120 | 180 | 240 | 300 | 600 | 900 | 1200 (seconds)

Function: Specifies whether the IP router times out entries in the address-resolution cache for this interface, and specifies the timeout interval in seconds if the interface does time out entries. The address-resolution cache contains host physical addresses learned by means of ARP or Proxy ARP. A host entry is timed out (deleted) if the IP router sends no traffic destined for that host within the specified timeout period.

Instructions: Select Off to disable timeout on this interface; the IP router does not time out address-resolution cache entries. Select one of the other options to enable timeout with a timeout interval equal to the value you select (for example, 120 seconds); the IP router removes address-resolution cache entries that have not been accessed within the specified number of seconds. Once an entry is removed, the IP router must use ARP to re-acquire the physical-level address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.15

Parameter: Checksum

Path: Configuration Manager > Protocols > IP > Interfaces

Default: On

Options: On | Off

Function: Specifies whether UDP checksum processing is enabled on this interface.

Instructions: Select On to enable UDP checksum processing for the interface; all outgoing and incoming UDP datagrams are subject to checksumming. You should select On in virtually all instances. Select Off to disable UDP checksum processing and provide backward compatibility with UNIX BSD 4.1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.16

Parameter: MAC Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: 0 | a user-specified MAC address | if the interface is on an SMDS circuit, the entire E.164 address -- for example, C1 617 555 5000 FFFF

Function: Specifies a MAC (media access control) address for this IP interface. The IP router will use its IP address and this MAC address when transmitting and receiving packets on this interface.

Instructions: Enter 0 to configure the IP router to use its IP address and the circuit's MAC address when transmitting packets on this interface. Enter your own MAC address to configure the IP router to use its IP address and the specified MAC address when transmitting packets on this interface. If the interface is on an SMDS circuit, by default, IP uses the SMDS-configured address. To configure this parameter for a multinet or multigroup configuration, refer to *Configuring SMDS*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.17

Parameter: TR Endstation

Path: Configuration Manager > Protocols > IP > Interfaces
Default: Off
Options: On | Off
Function: Specifies source routing over Token Ring selection.
Instructions: Use the On option to enable the parameter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.64

Parameter: Redirects

Path: Configuration Manager > Protocols > IP > Interfaces
Default: Enable
Options: Enable | Disable
Function: Indicates whether or not this interface sends out ICMP redirects. ICMP redirects are messages sent by the router to alert a host that it should be using a different path to route data.
Instructions: Reset to Disable if you do not want this interface to send out redirects. For example, in a Frame Relay network, two stations on the same network may not be directly connected if the network is not fully meshed. Thus, in this case, you would set Redirects to Disable.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.70

Parameter: **Enet Arp Encaps**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: ARP Ethernet

Options: ARP Ethernet | ARP SNAP | ARP Both | Probe LSAP | ARP Ethernet/Probe LSAP | ARP SNAP/Probe LSAP | ARP Both/Probe LSAP

Function: Defines the datalink encapsulation to use for ARP and HP Probe packets generated at this interface if the underlying medium is Ethernet. This parameter is ignored if the underlying medium is anything other than Ethernet.

Instructions: Depending on the selection you have made for the ARP Resolution parameter (ARP, Probe, or ARP/Probe), select the appropriate encapsulation option. If your address-resolution scheme is ARP only, select Ethernet encapsulation, SNAP encapsulation, or Ethernet/SNAP encapsulation. If your resolution scheme is HP Probe only, select LSAP encapsulation. If your resolution scheme is ARP/Probe, select Ethernet/LSAP encapsulation, SNAP/LSAP encapsulation, or Ethernet/SNAP/LSAP encapsulation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.71

Parameter: **SMDS Group Address**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: The SDMS-configured address

Options: A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider

Function: Provides a MAC-layer multicast address for this IP interface in an SMDS network. This parameter is displayed only if this is an SMDS circuit.

Instructions: Enter an entire E.164 address -- for example, E1 617 555 1212 FFFF. If you do not supply an address, IP uses the SDMS-configured address. To configure this parameter for a multinet or multigroup configuration, refer to *Configuring SMDS.*

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.65

Parameter: SMDS Arp Req Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: The SDMS-configured address

Options: A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider

Function: Provides an address-resolution multicast address for this IP interface in an SMDS network. This parameter is only displayed if this is an SMDS circuit.

Instructions: Enter an entire E.164 address -- for example, E1 617 555 1212 FFFF. If you do not supply an address, IP uses the SDMS-configured address. To configure this parameter for a multinet or multigroup configuration, refer to *Configuring SMDS*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.66

Parameter: WAN Broadcast

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a broadcast address for this IP interface in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will broadcast the message. This parameter is displayed only if this is a Frame Relay circuit.

Instructions: Enter the broadcast address provided by the Frame Relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.67

Parameter: WAN Multicast #1

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a multicast address for this IP interface that will send messages to all OSPF routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF routers. This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF routers as provided by the Frame Relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.68

Parameter: WAN Multicast #2

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a multicast address for this IP interface that will send messages to all OSPF designated routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF designated routers. This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF designated routers as provided by the Frame Relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.69

Parameter:	**Slot Mask**
Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Slot-mask bit set to 1 (enabling circuitless IP interface support) for every router slot running IP
Options:	For each slot in the router, Site Manager allows you to set the slot-mask bit to 1 (circuitless IP interface support enabled) or 0 (circuitless IP interface support disabled)
Function:	Specifies whether circuitless IP interface support is enabled or disabled on each slot in the router.
Instructions:	If you have configured a circuitless IP interface and do not wish it to run on certain slots, set the slot-mask bit to 0 on those slots. Be certain to keep the slot-mask bit set to 1 on at least one slot running IP; otherwise, the circuitless IP interface will not initialize. Setting the slot-mask bit parameter to 1 on an empty slot, a slot containing a system resource module, or a slot with no IP support does not affect the circuitless IP interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.1.75

Parameter:	**Max Forwarding Table Size**
Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	128 entries
Options:	64-entry minimum; no maximum
Function:	Specifies the maximum number of entries allowed in the forwarding table at one time.
Instructions:	Specify a forwarding table size for each interface. This parameter controls the number of destinations that are cached in the forwarding table on this receiving interface. When this interface receives an IP packet, the router looks up the destination in the forwarding table. Therefore, an interface that receives packets that are destined for a large number of different destinations may benefit from a larger forwarding table. The larger the number of entries, the more likely it is that the destination will already be in the forwarding table and the faster the route lookups will be for those destinations. Configuring a forwarding table size that is larger than necessary reduces the total amount of memory usable by other applications. Configuring a routing table too small can affect overall router performance. A check of the number of cache hits and misses will help determine the optimal size of the forwarding table. For debugging purposes, if you see the wfIpInterfaceCacheMisses statistic going up at an alarming rate, you should consider increasing the table size. However, an occasional cache miss does not warrant an increase in table size.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.1.104

Parameter: Enable Security

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disable

Options: Enable │ Disable

Function: Specifies whether Revised IP Security Option (RIPSO) is enabled for the interface.

Instructions: If you do not support RIPSO on your network, simply accept the default setting, Disable. If you are configuring RISPO support, set this parameter to Enable. Then see "Configuring RIPSO Support" for instructions on setting the rest of the RIPSO parameters that you must configure. Once you set this parameter to Enable, you can access the rest of the RIPSO parameters. If you do not enable this parameter, Site Manager does not activate the RIPSO parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.76

Parameter: UnNumbered Assoc Alternate

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disable

Options: Enable │ Disable

Function: Automatically assigns an alternate associated address to an unnumbered interface in the event that the primary associated address has gone down. IP uses the first available interface.

Instructions: Use the alternate unnumbered address option to ensure that the unnumbered interface has a usable associated address on the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.111

Parameter:	**ATM ARP Mode**
Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Client
Options:	Client | Server
Function:	Specifies whether the router is running as an ATM client or server on this interface.
Instructions:	You must configure one ATMARP server for each logical IP subnet you define.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.1.112

Parameter:	**ARP Server ATM Address Network Prefix**
Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	None
Options:	XX00000000000000000000000000 to XXFFFFFFFFFFFFFFFFFFFFFF where XX = 39, 45, or 47
Function:	Defines the ATM address network prefix of the ATMARP server on your network.
Instructions:	Enter the ATM address network prefix of the ATMARP server on your network. A complete ATM address consists of a network prefix and a user part. Use the ARP Server ATM Address User Part parameter to supply the user part of the ATM address.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.1.113

Parameter: **ARP Server ATM Address User Part**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: XX00000000000000 to FEFFFFFFFFFFFF

Function: Defines the user part (suffix) of the ATM address for the ATM ARP server on your network. The user part suffix consists of a 6-byte endstation identifier and a 1-byte selector field.

Instructions: Enter the user part suffix of the ATM ARP server on your network. A complete ATM address consists of a network prefix and a user part. Use the ARP Server ATM Address Network Prefix parameter to supply the network part of the ATM address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.113

Parameter: **Registration Refresh Interval**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 900 seconds for a client, 1200 seconds for a server

Options: Any interval (in seconds)

Function: For a client, this parameter specifies the interval between registration refreshes. For a server, this parameter specifies the duration for which the registration is valid.

Instructions: Determine whether ATMARP is running as a client or as a server on this interface and enter an appropriate value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.115

Parameter: **TR Endstation ARP Type**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: STE

Options: STE | ARE

Function: Specifies the ARP type for an interface configured for Token Ring support.

Instructions: For spanning tree explorer (STE) ARP packets, use the default. For all route explorer (ARE) packets, select ARE. Set the TR Endstation parameter to ON.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.127

IP Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Global

Default: This parameter defaults to Enable when you add IP support to a circuit.

Options: Enable | Disable

Function: Specifies the state of the IP router software.

Instructions: Select Enable if you have previously disabled the IP router software and now wish to re-enable it. Select Disable to disable the IP router software. In dynamic mode, when you set this parameter to Disable, you immediately prohibit all Site Manager communication with the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.2

Parameter: Forwarding

Path: Configuration Manager > Protocols > IP > Global

Default: Forwarding

Options: Forwarding | Not Forwarding

Function: Specifies whether the IP router forwards IP traffic that is not explicitly addressed to it.

Instructions: Select Forwarding if you want the IP router to route (forward) IP traffic. Forwarding configures the IP router to process all broadcast packets and all IP packets explicitly addressed to it, and to route all other IP packets. Select Not Forwarding if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces but also want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access. Not Forwarding configures the IP router to act as an IP host; it does not forward IP traffic, but it still processes packets explicitly addressed to it. In Not Forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated. Because the IP router does not forward IP traffic in Not Forwarding mode, you must configure the router to bridge IP traffic not explicitly addressed to it. You must configure the bridge for each circuit that conveys IP datagrams. The bridge will then forward all IP datagrams that are not explicitly addressed to the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.4

Parameter: ARP Forwarding

Path: Configuration Manager > Protocols > IP > Global

Default: Forwarding

Options: Forwarding | Not Forwarding

Function: Specifies how ARP should act in relation to IP's forwarding state. Note that Forwarding means IP is in forwarding mode. If this parameter is set to Forwarding, then ARP packets are either consumed (if destined for the router) or dropped. If this parameter is set to Not Forwarding, ARP packets are consumed, if destined for the router, or bridged onto remaining ARP interfaces.

Instructions: Always set this parameter the way you set the Forwarding parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.3

Parameter: Nonlocal ARP Source

Path: Configuration Manager > Protocols > IP > Global

Default: Drop

Options: Drop | Drop and Log

Function: Determines what happens when IP encounters an invalid ARP source address. If this parameter is set to Drop and Log, IP logs an invalid ARP source address when processing an ARP request. If this parameter is set to Drop, IP does not log the invalid ARP source address. In either case, IP drops the invalid ARP request.

Instructions: If you want to log the invalid ARP source address, set the parameter to Drop and Log. Otherwise, set the parameter to Drop.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.4

Parameter: **Nonlocal ARP Destination**

Path: Configuration Manager > Protocols > IP > Global

Default: Drop

Options: Drop | Accept

Function: Determines whether IP drops ARP requests in which the source and destination addresses are located in different networks or subnetworks. This parameter allows Proxy ARP to generate replies when the source and destination networks in the ARP request are different.

Instructions: To process ARP requests with source and destination addresses from different networks, set the parameter to Accept. The Proxy parameter must be set to Enable for the router to generate ARP replies.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.5

Parameter: **Default TTL**

Path: Configuration Manager > Protocols > IP > Global

Default: 30

Options: 1 to 255 hops

Function: Specifies the starting value of the Time to Live (TTL) counter for each packet the router originates and transmits (called a source packet). When the router transmits a source packet, the TTL counter starts to decrement. Each router, or hop, that the packet traverses decrements the TTL counter by one. When the counter reaches zero, the router discards the packet unless it is destined for a locally attached network. The TTL counter prevents packets from looping endlessly through the network.

Instructions: Enter the maximum number of hops a source packet can traverse.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.5

Parameter: RIP Diameter

Path: Configuration Manager > Protocols > IP > Global

Default: 15

Options: 1 to 127

Function: Specifies the value, or hop count, the Routing Information Protocol (RIP) uses to denote infinity. In order for RIP to operate properly, every router within the network must be configured with an identical RIP diameter value. If RIP is not enabled, this parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, the IP router still must understand network width.

Instructions: You must set this parameter so that none of the interface cost, static cost, or route filter cost parameters exceed the RIP diameter. We recommend that you accept the default RIP diameter value of 15.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.6

Parameter: Zero Subnet Enable

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable | Disable

Function: Specifies whether an interface address whose subnet portion is all zeros or all ones should be declared legal or not. If you set this parameter to Enable, then you can configure IP interfaces with a subnet ID of zero. Setting this parameter to Disable prevents you from doing so.

Instructions: Accept the default, Disable, if you do not have any interfaces that have a zero subnet ID. Otherwise, reset this parameter to Enable. The use of all-zero subnet addresses is discouraged for the following reason: if an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all-subnets broadcast from a directed broadcast for the zero subnet.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.10

Parameter: Estimated Networks

Path: Configuration Manager > Protocols > IP > Global

Default: 0

Options: 0 to 2147483647

Function: Allows the IP software to preallocate system resources based on the anticipated size of the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation. Preallocation also makes better use of memory and reduces the amount of memory required.

Instructions: Set to the number of networks (including unique subnets) that you expect. Avoid using a number that is excessively large. This will cause a wasteful overallocation of memory. If you use the default value, 0, IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.11

Parameter: Estimated Hosts

Path: Configuration Manager > Protocols > IP > Global

Default: 0

Options: 0 to 2147483647

Function: Allows the IP software to preallocate system resources based on the anticipated size of the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation.

Instructions: Set to the number of hosts that you expect. Avoid using a number that is excessively large. This will cause a wasteful overallocation of memory. If you use the default value, 0, IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.13

Parameter: Enable Default Route for Subnets

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable | Disable

Function: Specifies whether the IP router uses a default route for unknown subnets. The default route must be present in the routing table. When you set this parameter to Enable, the IP router uses a default route. When you set this parameter to Disable, the IP router does not use a default route.

Instructions: Accept the default, Disable, if you do not want the IP router to use a default route for unknown subnets. Otherwise, reset this parameter to Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.14

Parameter: Maximum Policy Rules

Path: Configuration Manager > Protocols > IP > Global

Default: 32

Options: Any integer

Function: Specifies the maximum number of policy rules that can be configured per policy type (Accept or Announce) per protocol.

Instructions: To configure more than 32 Accept or Announce policy rules for a protocol, you must set this parameter to a larger value. IP will round the value up to the next multiple of 32.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.15

Parameter: Route Filter Support

Path: Configuration Manager > Protocols > IP > Global

Default: Enable

Options: Enable | Disable

Function: Specifies whether or not IP supports route filters.

Instructions: If you do not require support for route filters, select Disable. Otherwise, use the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.16

Parameter: **RIP Maximum Equal Cost Paths**

Path: Configuration Manager > Protocols > IP > Global

Default: 1

Options: 1 to 5

Function: Specifies the maximum number of equal-cost paths allowed for a network installed in the routing table by RIP.

Instructions: Use the IP global Multipath Method parameter to enable multipath costs and specify the method that IP uses to choose the next hop for a datagram.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.17

Parameter: **Multiple Next Hop Calculation Method**

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable │ Disable

Function: Enables and disables equal-cost multipath support for RIP and OSPF and specifies the method that IP uses to choose the next hop when more than one is available. Three methods are available: round-robin selection, selection based on the source addresses (IP forwards all packets with the same source address to the same next hop), and selection based on the source and destination address (IP forwards all packets with the same source and destination address to the same next hop).

Instructions: Configure RIP and OSPF to support equal-cost routes to the same destination.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.18

Parameter: **ISP Mode**

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable │ Disable

Function: Enables and disables internet service provider (ISP) features.

Instructions: Use this parameter to configure BGP as a soloist and to disable the use of forwarding tables on IP interfaces.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.19

IP Accounting Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Global
Default: Enable
Options: Enable | Disable
Function: Enables and disables IP Accounting on the router.
Instructions: Use this parameter to disable and re-enable IP Accounting.
MIB Object ID: 1.3.6.1.4.1.18.3.5.20.1.1#2

Parameter: Threshold

Path: Configuration Manager > Protocols > IP > Global
Default: 512 entries
Options: 1 to 10,240 entries
Function: Specifies the maximum number of entries in the IP Accounting table.
Instructions: Specify a maximum number that meets the requirements of IP Accounting on this router.
MIB Object ID: 1.3.6.1.4.1.18.3.5.20.1.1#3

Parameter: Trap Percent

Path: Configuration Manager > Protocols > IP > Global
Default: 80 percent
Options: 1 to 100 percent
Function: Specifies a value (a percentage of the maximum number of entries in the accounting table) that causes IP Accounting to send a trap message.
Instructions: Specify a percentage that meets the requirements of IP Accounting on this router.
MIB Object ID: 1.3.6.1.4.1.18.3.5.20.1.1#5

Parameter:	**IP Accounting Checkpoint Flagt**
Path:	Configuration Manager > Protocols > IP > Global
Default:	0
Options:	0 to 0x7FFFFFFF
Function:	Allows you to specify when IP Accounting takes a snapshot of the active table iand puts it in the checkpoint table.
Instructions:	Specify a flag value that meets the requirements of IP Accounting on this router.
MIB Object ID:	1.3.6.1.4.1.18.3.5.20.1.1#7

Static Route Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > Static Routes
Default:	This parameter defaults to Enable when you configure the static route.
Options:	Enable | Disable
Function:	Specifies the state (active or inactive) of the static route record in the IP routing tables.
Instructions:	Select Disable to make the static route record inactive in the IP routing table; the IP router will not consider this static route. Select Enable to make the static route record active again in the IP routing table.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.5.1.2

Parameter: Destination IP Address

Path: Configuration Manager > Protocols > IP > Static Routes

Default: None

Options: Any valid IP network address

Function: Specifies the IP address of the network to which you want to configure the static route. Specifies a supernet for which you want to configure a black hole static route.

Instructions: Enter the destination IP address in dotted decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet address. You can configure up to 12 static routes to the same destination.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.3

Parameter: Address Mask

Path: Configuration Manager > Protocols > IP > Static Routes

Default: None

Options: Based on the network class of the IP address you specified at the Destination IP Address parameter

Function: Specifies the subnet mask of the destination network. Specifies the supernet mask of the supernet for which you want to configure a black hole static route.

Instructions: Enter the subnet or supernet mask in dotted decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet mask.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.4

Parameter: Cost

Path: Configuration Manager > Protocols > IP > Static Routes
Default: 1
Options: 1 to the value of the RIP Diameter parameter (maximum 126)
Function: Specifies the number of router hops a datagram can traverse before reaching the destination IP address. The IP router uses the cost value when determining the best route for a datagram to follow.
Instructions: Enter the number of router hops.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.5

Parameter: Next Hop Addr

Path: Configuration Manager > Protocols > IP > Static Routes
Default: 0.0.0.0
Options: Any valid IP address
Function: Specifies the IP address of the next-hop router. Defines a black hole route for a supernet.
Instructions: Enter the IP address in dotted decimal notation. To configure a black hole static route, enter 255.255.255.255. If you are configuring a static route to an unnumbered interface, enter 0.0.0.0.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.6

Parameter: Next Hop Mask

Path: Configuration Manager > Protocols > IP > Static Routes
Default: 0.0.0.0
Options: Any valid subnet mask address
Function: Specifies the subnet mask of the next hop router. The parameter also defines a black hole route for a supernet.
Instructions: Enter the subnet mask in dotted decimal notation. To configure a black hole static route, enter 255.255.255.255. If you are configuring a static route to an unnumbered interface, enter 0.0.0.0.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.7

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Static Routes

Default: 16

Options: 1 to 16

Function: Specifies a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination.

Instructions: Enter a value from 1 to 16 for this static route. To configure a black hole static route, enter the maximum preference value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.8

Parameter: Unnumbered CCT Name

Path: Configuration Manager > Protocols > IP > Static Routes

Default: None

Options: A valid circuit name

Function: Specifies the local router circuit associated with the static route over an unnumbered interface.

Instructions: An entry for a route using an unnumbered interface must include the circuit associated with the interface. Use this parameter to specify that circuit name.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.11

Adjacent Host Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: Enable
Options: Enable | Disable
Function: Specifies the state (active or inactive) of the adjacent host in the IP routing tables.
Instructions: Select Disable to make the adjacent host record inactive in the IP routing table; the IP router will not consider this adjacent host. Select Enable to make the adjacent host record active again in the IP routing table.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.2

Parameter: IP Adjacent Host Address

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: None
Options: Any valid IP address
Function: Specifies the IP address of the device for which you want to configure an adjacent host.
Instructions: Enter the IP address in dotted decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.3

Parameter: Next Hop Interface Addr

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: 0.0.0.0
Options: A valid IP address
Function: Specifies the IP address of the router's network interface to the adjacent host.
Instructions: Enter the IP address in dotted decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.4

Parameter: Next Hop Interface Mask

Path: Configuration Manager > Protocols > IP > Adjacent Hosts

Default: 0.0.0.0

Options: Based on the network class of the IP address specified at the Next Hop Interface Addr parameter

Function: Specifies the subnet mask of the IP address specified for the Next Hop Addr parameter.

Instructions: Enter the subnet mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.5

Parameter: MAC Address

Path: Configuration Manager > Protocols > IP > Adjacent Hosts

Default: None

Options: Depend on the data link you have selected

Function: Specifies the physical address of the adjacent host. This value can be a 48-bit Ethernet address, a 64-bit SMDS address, an ATM PVC VPI/VCI address, or, for an ATM SVC, the address of the ATM interface.

Instructions: Enter the MAC address as a 12-digit hexadecimal number. Enter an ATM/PVC address in the form Virtual Path Identifier/Virtual Channel Identifier -- for example, 0/32.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.6

Parameter: Host Encapsulation

Path: Configuration Manager > Protocols > IP > Adjacent Hosts

Default: Ethernet

Options: Ethernet | SNAP | PDN | DDN | SNAPIP | NULL

Function: Specifies the adjacent host's encapsulation method.

Instructions: Select Ethernet or SNAP (Service Network Access Point) if you are defining a point-to-point network interface or if the adjacent host resides on an Ethernet. For an X.25 interface, select PDN or DDN. For an adjacent host on an ATM logical IP subnet, select SNAP. (SNAPIP and NULL also specify host encapsulation methods for ATM networks.)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.7

Parameter: **Adjacent Host X.121 Address**

Path: Configuration Manager > Protocols > IP > Adjacent Hosts

Default: None

Options: Any valid X.121 address

Function: Specifies the X.121 address of the adjacent host. Set this parameter only if this is a PDN/X.25, DDN/X.25, or BFE/X.25 connection.

Instructions: Enter the appropriate X.121 address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.9

TFTP Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > TFTP

Default: Enable

Options: Enable | Disable

Function: Specifies whether TFTP is enabled for the IP router.

Instructions: Select Enable to enable TFTP for the IP router. Because TFTP allows write-access to the router's file system, we recommend that you do not enable TFTP in network environments in which you are concerned with security. Select Disable to disable TFTP for the IP router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.1

Parameter: **Default Volume**

Path: Configuration Manager > Protocols > IP > TFTP

Default: 2

Options: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14

Function: Specifies which of the router's slots will be used, by default, for all TFTP GETs and PUTs.

Instructions: Specify the appropriate slot number. If you are configuring an AN™ router, you must specify slot 1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.2

Parameter: Retry Time Out

Path: Configuration Manager > Protocols > IP > TFTP
Default: 5 seconds
Options: Any number of seconds
Function: Specifies the number of seconds TFTP waits for an acknowledgment before retransmitting the last packet.
Instructions: Specify the appropriate number of seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.4

Parameter: Close Time Out

Path: Configuration Manager > Protocols > IP > TFTP
Default: 25 seconds
Options: Any number of seconds
Function: Specifies the number of seconds TFTP waits, after it has successfully received a file, to make sure that the sender has received the last acknowledgment.
Instructions: Specify the appropriate number of seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.5

Parameter: Retransmit

Path: Configuration Manager > Protocols > IP > TFTP
Default: 5 retransmissions
Options: Any number of retransmissions
Function: Specifies the number of times TFTP retransmits an unacknowledged message before abandoning the transfer attempt.
Instructions: Specify the number of retransmissions.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.6

RIPSO Parameters

Parameter: Enable Security

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enable

Options: Enable | Disable

Function: Enables or disables IP security options for this interface.

Instructions: Set to Disable if you want to disable IP security options. If you set this parameter to Disable, then the router accepts only the following IP datagrams: labeled IP datagrams with the classification level set to Unclassified and no authority flags set and unlabeled IP datagrams.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.76

Parameter: Strip Security

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: None | Incoming | Outgoing | All

Function: Specifies the type of IP datagram from which the router should remove the IP security options.

Instructions: Select the type of IP datagram from which you want IP security options to be removed. None causes the router to leave IP security options on all inbound and outbound IP datagrams intact. Incoming causes the router to strip the IP security option from each incoming IP datagram, after checking the IP datagram against the interface's security configuration. Outgoing causes the router to strip the IP security option from each outgoing IP datagram, before checking each datagram against the interface's security configuration. All causes the router to strip the IP security options from both incoming and outgoing IP datagrams: incoming datagrams after checking each against this interface's security configuration and outgoing datagrams before checking each against the interface's security configuration. If you set this parameter to Outgoing or All, then you must set the Require Out Security parameter to None. (Similarly, if you set the Require Out Security parameter to Forwarded, Originated, or All, then you must set this parameter to None or Incoming.)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.77

Parameter: Require Out Security

Path: Configuration Manager > Protocols > IP > Interfaces

Default: All

Options: None │ Forwarded │ Originated │ All

Function: Specifies which type of outbound datagrams require IP security labels.

Instructions: Select None: the router forwards unlabeled IP datagrams unchanged on this interface. In addition, those IP datagrams that it originates and transmits do not require labels. Select Forwarded: the router requires all IP datagrams it forwards on this interface (not those it originates) to contain basic IP security options. If the datagram already contains an IP security label, the router forwards the datagram unchanged. If the datagram is unlabeled, the router adds the implicit or default label to the datagram before forwarding it. Select Originated: the router specifies basic IP security options for all IP datagrams it originates and transmits on this interface. The router adds the default label to IP datagrams it originates and transmits on this interface. Select All: the router requires all datagrams (both those that it forwards and those it originates) on this interface to contain basic IP security options. It supplies the implicit or default label for those datagrams that do not already contain one. If you set this parameter to Originated or All, then you must enable the Default Label and Error Label parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.78

Parameter: Require In Security

Path: Configuration Manager > Protocols > IP > Interfaces

Default: All

Options: None │ All

Function: Specifies which type of incoming IP datagram requires security labels.

Instructions: Select None: the router does not require inbound IP datagrams to contain labels. Select All: the router requires all inbound IP datagrams received on this interface to contain basic IP security options.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.79

Parameter: Min Level

Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Unclassified
Options:	Unclassified | Confidential | Secret | Top Secret
Function:	Specifies the minimum security level that the router allows for inbound or outbound IP datagrams. This parameter, together with the Max Level parameter, specifies the range of classification levels that the router will accept and process. The router drops IP datagrams it receives on this interface that are below the minimum level specified here.
Instructions:	Select a minimum security level for this interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.80

Parameter: Max Level

Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Top Secret
Options:	Unclassified | Confidential | Secret | Top Secret
Function:	Specifies the maximum security level that the router allows for inbound or outbound IP datagrams. This parameter, together with the Min Level parameter, specifies the range of classification levels that the router accepts. The router drops IP datagrams it receives or transmits on this interface that are above the maximum level specified here.
Instructions:	Select a maximum security level for this interface. The maximum level must be greater than or equal to the minimum level.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.81

Parameter: Must Out Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies which authority flags must be set in the protection authority field of all outbound datagrams.

Instructions: Select all of those authority flags that the router must set in all outbound IP datagrams it transmits on this interface. If you do not select any authority flags (the default setting), the router does not set any protection authority flags in outbound IP datagrams.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.82

Parameter: May Out Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: ANY

Options: ANY | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies which authority flags may be set in the protection authority field of all outbound datagrams. The authorities you specify here must be a superset of the authorities you specify for the Must Out Authority parameter.

Instructions: The default setting specifies that any of the authority flags may be set. Either accept the default setting or reset and select only those authority flags that are appropriate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.83

Parameter: Must In Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies which authority flags must be set in the protection authority field of inbound IP datagrams.

Instructions: Select all of those authority flags that must be set in inbound IP datagrams received on this interface. If you do not select any authority flags (the default setting), then the router does not require a datagram to have authority flags set, but still accepts the datagram if any flags are set.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.84

Parameter: May In Authorit

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Any

Options: Any | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies which authority flags may be set in the protection authority field of inbound IP datagrams. The authorities you specify here must be a superset of the authorities you specify for the Must In Authority parameter.

Instructions: The default setting specifies that any of the authority flags may be set. Either accept the default setting or reset and select only those authority flags that are appropriate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.85

Parameter: Implicit Label

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enable

Options: Enable | Disable

Function: If you select Enable, the router uses the Implicit Authority and Implicit Level fields to create an implicit label. The router supplies the implicit label to unlabeled inbound datagrams received by this interface. If you select Disable, the router does not supply implicit labels for this interface.

Instructions: Accept the default, Enable, to allow the router to supply implicit labels for unlabeled inbound datagrams.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.86

Parameter: Implicit Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies the authority flags that the router sets when it supplies implicit security labels for unlabeled inbound IP datagrams.

Instructions: Select all of those authority flags that the router should set when it supplies an implicit security label. The set of authority flags you specify here must include the set of authority flags you specified for the Must In Authority parameter, and cannot include any of the flags you did not specify for the May In Authority parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.87

Parameter: Implicit Level

Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Unclassified
Options:	Unclassified | Confidential | Secret | Top Secret
Function:	Specifies the security level that the router sets when it supplies implicit security labels for unlabeled, inbound IP datagrams.
Instructions:	Specify a level within the range specified by the Min Level and Max Level parameters.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.88

Parameter: Default Label

Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Enable
Options:	Enable | Disable
Function:	If you select Enable, the router uses the Default Authority and Default Level fields to create a default label. The router supplies the default label to unlabeled outbound datagrams originated or forwarded out this interface. If you select Disable, the router does not supply default labels for this interface.
Instructions:	To allow the router to supply default labels for unlabeled outbound datagrams, accept the default, Enable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.89

Parameter: Default Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESI | SCI | NSA | DOE

Function: Specifies the authority flags that the router uses when it supplies default security labels to unlabeled outbound IP datagrams.

Instructions: Select those authority flags that the router should set when it supplies default security labels. The set of authority flags you specify must include the set of authority flags specified for the Must Out Authority parameter, and cannot include any of the flags you did not specify for the May Out Authority parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.90

Parameter: Default Level

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Unclassified

Options: Unclassified | Confidential | Secret | Top Secret

Function: Specifies the security level that the router sets when it supplies default security labels to unlabeled outbound IP datagrams.

Instructions: Specify a default level within the range specified by the Min Level and Max Level parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.91

Parameter: Error Label

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enable

Options: Enable | Disable

Function: If you select Enable, the router uses the Error Authority and Min Level fields to create an error label. The router supplies the error label to outbound ICMP error datagrams. If you select Disable, the router does not supply error labels for this interface.

Instructions: To allow the router to supply error labels for outbound ICMP error datagrams, accept the default, Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.92

Parameter: Error Authority

Path: Configuration Manager > Protocols > IP > Interfaces

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESI | SCI | NSA | DOE | ALL

Function: Specifies the authority flags that the router uses when it supplies error security labels to outbound ICMP error datagrams.

Instructions: Select those authority flags that the router should set when it supplies error security labels to outbound ICMP error datagrams. The set of authority flags you specify here must include the set of authority flags you specified for the Must Out Authority parameter, and cannot include any of the flags you did not specify for the May Out Authority parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.93

Router Discovery Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: Enable
Options: Enable | Disable
Function: Disables and enables Router Discovery on this interface.
Instructions: If you configured this interface with Router Discovery, use this parameter to disable Router Discovery.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.2

Parameter: Broadcast Type

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: Multicast
Options: Multicast | Local | Direct
Function: Specifies the type of broadcast to use in sending advertisements.
Instructions: Use Multicast wherever possible; that is, on any link where all listening hosts support IP multicast.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.5

Parameter: Minimum Interval

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 450
Options: A value specifying the number of seconds
Function: Specifies the minimum time interval between advertisements.
Instructions: Specify a value that is no less than 3 seconds and less than the value you set for the Maximum Interval parameter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.6

Parameter: Maximum Interval

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 600
Options: A value specifying the number of seconds
Function: Specifies the maximum time interval between advertisements.
Instructions: Specify a value that is not less than 4 seconds, is greater than the value you specified for the Minimum Interval parameter, and is not greater than 1800 seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.7

Parameter: Lifetime

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 1800
Options: A value specifying the number of seconds
Function: Specifies the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements.
Instructions: Specify a value that is no less than the value you set for the Maximum Interval parameter and no greater than 9000 seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.8

Parameter: Interface Preference

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 0
Options: A numeric value
Function: Specifies the preferability (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet.
Instructions: Enter a value indicating the relative preferability of the router address. Enter a preference value of 0x80000000 to indicate to neighboring hosts that the address is not to be used as a default route.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.9

NetBIOS

NetBIOS Global Parameters

Parameter: **Enable/Disable**
Path: Configuration Manager > Protocols > IP > NetBIOS > Global
Default: Enable
Options: Enable │ Disable
Function: Enables or disables NetBIOS on this router.
Instructions: If NetBIOS has been configured on this router, use this parameter to disable and re-enable it as required.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.2

Parameter: **NetBIOS Name Caching**
Path: Configuration Manager > Protocols > IP > NetBIOS > Global
Default: Disable
Options: Enable │ Disable
Function: Globally enables or disables the ability of the router to cache the name associated with each NetBIOS server that is active on the network.
Instructions: Select Enable to activate NetBIOS server name caching at every NetBIOS interface configured on the node.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.4

Parameter: **15-Character NetBIOS Name Caching**
Path: Configuration Manager > Protocols > IP > NetBIOS > Global
Default: Disable
Options: Enable │ Disable
Function: Enables or disables the ability of the router to treat a NetBIOS name as either a 15- or a 16-character entity.
Instructions: Select Enable to activate 15-character NetBIOS name caching at every NetBIOS interface configured on this router. Select Disable if you want NetBIOS to treat names as 16-character entities.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.5

Parameter: **Create MIB Inst for Cached Name**

Path: Configuration Manager > Protocols > IP > NetBIOS > Global

Default: Enable

Options: Enable | Disable

Function: Enables or disables the ability of the system to

Instructions: Create a MIB instance for each name entry stored in the name cache. Delete a MIB instance for each NetBIOS name entry that ages out of the name cache. Select Disable if you want to release the system memory and processing resources otherwise dedicated to maintaining cached names in the MIB.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.6

Parameter: **Max Name Cache Entries**

Path: Configuration Manager > Protocols > IP > NetBIOS > Global

Default: 100 entries

Options: 1 to 2147483647 entries

Function: Specifies the maximum number of entries you need to provide in the NetBIOS name cache.

Instructions: You can adjust the value of this parameter in direct proportion to the total number of server names expected to be active during intervals of peak traffic load or performance demand on the router. A value of 100 is suitable for networks that include up to 100 NetBIOS names to cache.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.7

Parameter: Name Cache Age

Path: Configuration Manager > Protocols > IP > NetBIOS > Global

Default: 300 seconds

Options: Any value that can rapidly age infrequently referenced names out of the NetBIOS name cache

Function: Specifies an age (in seconds) when inactive NetBIOS names expire from the NetBIOS name cache.

Instructions: Choose an aging value that allows infrequently referenced or obsolete server names to expire from the name cache. The smaller the value, the less efficient broadcast reduction is, but the more quickly the network recovers topology changes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.9

Parameter: Hash Entry Count

Path: Configuration Manager > Protocols > IP > NetBIOS > Global

Default: 253

Options: Any integer value

Function: Specifies the number of entries you want to allow in the cache lookup tables. Each NetBIOS interface has a local table to store and retrieve the names of NetBIOS servers active on the network.

Instructions: For networks that actively use up to 2500 NetBIOS server names, use the default value (253). To determine a hash entry count for larger networks, divide the total number of unique NetBIOS server names active in the network by 10; adjust the quotient to the nearest (higher or lower) prime number; and replace the default value with the new, calculated number. Increasing the number of hash table entries does not increase the number of names that a router can cache. With larger networks, increasing the size of the hash tables may, however, reduce internal cache lookup time, thereby improving overall performance.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.10

Parameter:	**Rebroadcast Packet TTL**
Path:	Configuration Manager > Protocols > IP > NetBIOS > Global
Default:	5 seconds
Options:	1 to 255 seconds
Function:	Specifies the time-to-live value in seconds to use in rebroadcast packets.
Instructions:	Use this parameter to restrict the number of routers a rebroadcast packet can traverse. To prevent NetBIOS broadcast packets from traversing the network indefinitely, set the parameter to a minimal value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.1.13

Parameter:	**Rebroadcast Record Route**
Path:	Configuration Manager > Protocols > IP > NetBIOS > Global
Default:	Disable
Options:	Enable | Disable
Function:	Enables and disables the Insertion of Record Route option in rebroadcast packets.
Instructions:	If all IP entities support this option, select Enable to allow the NetBIOS entity in the router to determine whether it has received this packet before on this interface. If so, the router drops it. This option prevents rebroadcast packets from looping forever.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.1.14

NetBIOS/IP Interface Table Parameters

Parameter:	**Enable/Disable**
Path:	Configuration Manager > Protocols > IP > NetBIOS >Interface
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables NetBIOS on this IP interface.
Instructions:	If NetBIOS has been configured and enabled on the router, use this parameter to disable and re-enable it on this interface as required.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.2.1.2

Parameter: **NetBIOS Name Caching**

Path: Configuration Manager > Protocols > IP > NetBIOS >Interface

Default: Enable

Options: Enable │ Disable

Function: Enables or disables the ability of this interface to cache the name for each NetBIOS server active in the network.

Instructions: Select Enable if you disabled server name caching previously and you want now to re-enable that function. Select Disable if you want to release system memory and processing resources otherwise dedicated to server name caching.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.2.1.8

Parameter: **Enable NetBIOS Inbound Broadcasts**

Path: Configuration Manager > Protocols > IP > NetBIOS >Interface

Default: Enabled

Options: Enable │ Disable

Function: Enables or disables inbound broadcasts on this interface.

Instructions: If NetBIOS is configured and enabled on the router and enabled on this interface, use this parameter to enable and disable inbound broadcasts as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.2.1.9

Parameter: **Enable NetBIOS Outbound Broadcasts**

Path: Configuration Manager > Protocols > IP > NetBIOS >Interface

Default: Enable

Options: Enable │ Disable

Function: Enables or disables outbound broadcasts on this interface.

Instructions: If NetBIOS is configured and enabled on the router and enabled on this interface, use this parameter to enable and disable outbound broadcasts as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.2.1.10

Parameter:	**Rebroadcast Address**
Path:	Configuration Manager > Protocols > IP > NetBIOS >Interface
Default:	Null
Options:	An IP broadcast address
Function:	Specifies a broadcast address to use when rebroadcasting NetBIOS packets out this interface.
Instructions:	By default, NetBIOS uses the IP broadcast address configured for this interface. Set this parameter if you want to override this broadcast address.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.2.1.11

Static Entry Table Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > NetBIOS >Static Name
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables caching of the NetBIOS name you have selected.
Instructions:	Set the parameter to Enable to activate caching of the name you selected. Set the parameter to Disable to deactivate caching of the name you selected.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.4.1.2

Parameter:	**NetBIOS Scope ID**
Path:	Configuration Manager > Protocols > IP > NetBIOS >Static Name
Default:	None
Options:	A NetBIOS scope identifier
Function:	Identifies the area of the network across which the NetBIOS name is known.
Instructions:	Enter a name string that meets the requirements of the Domain Name System as described in RFC 833.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.11.4.1.5

Parameter: NetBIOS Station Name

Path: Configuration Manager > Protocols > IP > NetBIOS >Static Name

Default: None

Options: A name string of up to 16 characters

Function: Specifies the name of a NetBIOS station.

Instructions: Enter the NetBIOS name you want to add. The name must not exceed 16 characters. The system pads names shorter than 16 characters with ASCII space characters. To enter non-ASCII values in the name, use the form \x*bb*, where *bb* can be any two hexadecimal digits.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.4

Parameter: NetBIOS Scope ID

Path: Configuration Manager > Protocols > IP > NetBIOS >Static Name

Default: None

Options: A NetBIOS scope identifier

Function: Identifies the area of the network across which the NetBIOS name is known.

Instructions: Enter a name string that meets the requirements of the Domain Name System as described in RFC 833.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.5

Parameter: IP Address

Path: Configuration Manager > Protocols > IP > NetBIOS >Static Name

Default: None

Options: The IP address of the NetBIOS station

Function: Specifies an IP address to associate with the statically configured name.

Instructions: Enter a valid IP address of a NetBIOS station.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.6

OSPF Parameters

OSPF Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: Enable
Options: Enable | Disable
Function: Globally enables or disables OSPF on all router interfaces.
Instructions: Set to Disable if you want to disable OSPF for the entire router. Set to Enable if you previously disabled OSPF on the router and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.2

Parameter: Router ID

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: The IP address of the first OSPF circuit configured on this router.
Options: Any IP address; preferably, one of the router's IP interface addresses
Function: This IP address uniquely identifies this router in the OSPF domain. By convention, and to ensure uniqueness, one of the router's IP interface addresses should be used as the router ID. The router ID will determine the designated router on a broadcast link if the priority values of the routers being considered are equal. The higher the router ID, the greater its priority.
Instructions: Enter the appropriate IP address in dotted decimal notation. If both OSPF and BGP are running on the router, the OSPF router ID must be identical to the BGP identifier. In addition, the OSPF router ID must match one of the IP addresses configured on the router.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.4

Parameter: AS Boundary Router

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: No

Options: Yes | No

Function: Indicates whether or not this router functions as an AS boundary router. Only AS boundary routers are allowed to convert non-OSPF routes into OSPF routes so that they can be passed along throughout the OSPF routing domain. The router can be an AS boundary router if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

Instructions: Set this parameter to Yes if this router functions as an AS boundary router. Otherwise, accept the default value, No.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.7

Parameter: Hold Down Timer

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: 1 second

Options: 0 to 10 seconds

Function: Prevents the algorithm to compute a route from running more than once per holddown time. Its purpose is to free up the CPU. Note that a value of 0 means there is no holddown time.

Instructions: Either accept the default value of 1 second or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.9

Parameter: OSPF Slot

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: All slots

Options: Any slot on the router

Function: Indicates which slot(s) the OSPF soloist is eligible to run on. If the slot on which the OSPF soloist is running goes down, the router will attempt to run OSPF on another slot specified by this parameter.

Instructions: Select all of the appropriate slots. Use caution when selecting the slot(s) on which OSPF may run. If you choose an empty slot, and it is the only slot you choose, OSPF will not run; if you choose a slot that becomes disabled, and it is the only slot you choose, OSPF will not restart.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.10

Parameter: ASE Metric Support

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: Disable

Options: Enable | Disable

Function: Causes the router to use the route weight as the OSPF metric in OSPF ASE Type 2 advertisements.

Instructions: Disable ASE metric support if the router is to interoperate with routers using a pre-8.00 OSPF version. The new metric is not compatible with the pre-8.00 metric.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.11

Parameter: Backup Enable

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: Disable

Options: Enable | Disable

Function: Enables or disables the backup OSPF soloist's backup link state database. When the parameter is set to Disable, the OSPF backup soloist will not maintain a copy of the OSPF link state database.

Instructions: Select the default, Disable, if you do not want to back up the OSPF soloist.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.12

Parameter: Primary Log Mask

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: TRACE | INFO | DEBUG | INTF STATE | NBR STATE | BAD LS

Options: TRACE | INFO | DEBUG | INTF STATE | NBR STATE | LSA SELF-ORIGIN | LSA RECEIPT | ROUTE CHANGE | BAD LS | LESS RECENT LSA | MORE RECENT LSA | MAX AGE LSA

Function: Specifies which OSPF log messages should be logged in the primary log.

Instructions: Highlight the line entry for Primary Log Mask in the Edit OSPF Global Parameters window and click on Values. Choose the log messages that you want to enter into the primary log by clicking on their buttons. Then click on OK.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.13

Parameter: Backup Log Mask

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: Log no messages

Options: TRACE | INFO | DEBUG | INTF STATE | NBR STATE | LSA SELF-ORIGIN | LSA RECEIPT | ROUTE CHANGE | BAD LS | LESS RECENT LSA | MORE RECENT LSA | MAX AGE LSA

Function: Specifies which OSPF log messages should be logged in the backup log.

Instructions: Highlight the line entry for Backup Log Mask in the Edit OSPF Global Parameters window and click on Values. Choose the log messages that you wish to enter into the backup log by clicking on their buttons. Then click on OK.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.14

Parameter: **Tag Generation Method**

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: Zero

Options: Zero | Autotag | Proprietary

Function: Specifies the method of OSPF external tag field generation.

Instructions: Set the parameter to Autotag if you want OSPF to generate a tag value according to RFC 1403, “OSPF/BGP Interaction.” Use the default to insert 0 into the tag field. The Proprietary option is reserved for debugging purposes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.15

Parameter: **OSPF Maxium Path**

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: 1 path

Options: 1 to 5 equal-cost paths

Function: Specifies the maximum number of equal-cost paths allowed for a network installed by OSPF.

Instructions: If you have enabled equal-cost multpath support on the router, specify a value from 2 to 5.

MIB Object ID:

OSPF Interface Paramters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: Enable

Options: Enable │ Disable

Function: This parameter indicates whether or not OSPF is enabled on this interface. The default value, Enable, indicates that neighbor relationships may be formed on this interface, and that this interface will be advertised as an internal route to some area. The value Disable indicates that this is not an OSPF interface.

Instructions: Set this parameter to Disable if you do not want OSPF enabled on the interface. Set it to Enable if you previously disabled OSPF on this interface and now wish to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.2

Parameter: Area ID

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: 0.0.0.0

Options: Any 4-octet number in dotted decimal notation

Function: This parameter identifies the area to which this interface belongs.

Instructions: Enter the appropriate area ID in dotted decimal notation. Area ID 0.0.0.0 is reserved for the backbone.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.6

Parameter: Type

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: Broadcast

Options: Broadcast | NBMA (nonbroadcast multiaccess) | Point-to-point | Point-to-multipoint (STD) | Point-to-multipoint

Function: Indicates the type of network to which this interface is attached. Select Broadcast if this network is a broadcast LAN, such as Ethernet. Select NBMA if the network is a nonbroadcast network, such as X.25. Select Point-to-point for a synchronous, point-to-point interface. Select Point-to-multipoint (STD) if the network is a point-to-multipoint network. If you want to use the Bay Networks proprietary solution for Frame Relay point-to-multipoint networks, select Point-to-multipoint.

Instructions: Set this parameter to match this interface type. If you set this parameter to NBMA, you need to configure neighbors manually.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.7

Parameter: Rtr Priority

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: 1

Options: 0 to 255

Function: Indicates the priority of this interface. The router priority value is used in multiaccess networks (Broadcast, NBMA, or point-to-multipoint), for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network. In the case of equal Rtr Priority values, the router ID will determine which router will become the designated router. However, if there already is a designated router on the network when you start this router, it will remain the designated router no matter what your priority or router ID.

Instructions: Set the router priority to a value between 0 and 255 or accept the default value, 1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.8

Parameter: Transit Delay

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: 1 second

Options: 1 to 3600 seconds

Function: Indicates the estimated number of seconds it takes to route a packet over this interface.

Instructions: Either accept the default value of 1 second or enter some slightly higher number for slower-speed serial lines, for example, 15 to 20 seconds for a 19.8-KB line.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.9

Parameter: Retransmit Interval

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: 5 seconds

Options: 1 to 3600 seconds

Function: Indicates the number of seconds between link-state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting OSPF packets. Although the default value is 5, we suggest the following values for this parameter: for broadcast, 5 seconds; for point-to-point, 10 seconds; for NBMA, 10 seconds; for point-to-multipoint, 10 seconds.

Instructions: Either accept the default value of 5 seconds or set the retransmit interval to some slightly higher number for slower-speed serial lines.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.10

Parameter: Hello Interval

Path:	Configuration Manager > Protocols > IP > OSPF > Interfaces
Default:	10 seconds
Options:	1 to 65,535 seconds
Function:	Indicates the number of seconds between the hello packets that the router sends on the interface. Although the default value is 10 seconds, we suggest the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 20 seconds; for point-to-multipoint, 15 seconds.
Instructions:	Either accept the default value of 10 seconds or set the hello interval to some higher number for slower-speed serial lines. This value must be the same for all routers attached to the same network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.11

Parameter: Dead Interval

Path:	Configuration Manager > Protocols > IP > OSPF > Interfaces
Default:	40 seconds
Options:	1 to 2147483647 seconds
Function:	Indicates the number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. The dead interval value should be some multiple of the hello interval value. We suggest the following values for this parameter: for broadcast, 40 seconds; for point-to-point, 60 seconds; for NBMA, 80 seconds; for point-to-multipoint, 60 seconds.
Instructions:	Either accept the default value of 40 seconds or set the dead interval to some higher number for slower-speed serial lines. This value must be the same for all routers attached to the same network
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.12

Parameter:	**Poll Interval**
Path:	Configuration Manager > Protocols > IP > OSPF > Interfaces
Default:	120 seconds
Options:	1 to 2147483647 seconds
Function:	Indicates the largest number of seconds allowed between hello packets sent to an inactive nonbroadcast multiaccess neighbor.
Instructions:	Either accept the default value of 120 seconds or set this parameter to some slightly higher number for slower-speed serial lines.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.13

Parameter:	**Metric Cost**
Path:	Configuration Manager > Protocols > IP > OSPF > Interfaces
Default:	1
Options:	1 to 65535
Function:	Indicates the cost of using this type of service on this interface. We suggest the following values for this parameter: for > = 100 Mb/s, 1; for Ethernet/802.3, 10; for E1, 48/ for T1, 65; for 64 Kb/s, 1562; for 56 Kb/s, 1785, for 19.2 Kb/s, 5208; for 9.6 Kb/s, 10416. This parameter allows you to configure preferred paths. If you do want to configure a preferred path, allow that path to retain the default value of 1 or assign it a relatively low metric cost. Then, assign the less preferred paths a higher metric cost value.
Instructions:	Either accept the default value, 1, or enter a larger number for a slower path or a backup route.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.16

Parameter: Password

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: None

Options: Any ASCII string up to eight characters long

Function: Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. This parameter is valid only when Authentication Type is set to Simplepassword.

Instructions: Enter the appropriate password. All routers in the same area must either have no authentication or have the same password.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.17

Parameter: MTU Size

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: 1

Options: 1 | 2 | a number up to 10,000

Function: Specifies the maximum transmission unit (MTU) size of OSPF updates on this interface.

Instructions: Accept the default value, 1, to use the IP MTU size for that physical interface. Enter 2 to send packets no larger than the IP MTU size for Ethernet (1500). Enter a number up to 10,000 to specify an MTU size directly; the number you enter must be less than the IP MTU size for that physical interface. When running OSPF over a synchronous/PPP link, set the MTU size to a value less than the sync MTU size (1200). This allows all OSPF routes to be learned over the link.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.29

Neighbor Parameters for an NBMA Interface

Parameter: **Neighbor's IP Address**

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces > **Neighbors.** > **Add**.

Default: None

Options: IP address of neighbor

Function: Indicates by IP address a nonbroadcast multiaccess neighbor for this interface.

Instructions: Enter the appropriate IP address of the nonbroadcast multiaccess neighbor in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.4

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces > **Neighbors**.

Default: Enable

Options: Enable | Disable

Function: Allows you to enable and disable this neighbor configuration for this interface. This parameter is useful if you want to temporarily disable a neighbor configuration rather than delete it.

Instructions: Set to Disable if you want to disable this neighbor configuration. Or set to Enable if you previously disabled this neighbor configuration and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2

Parameter: Priority

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces > **Neighbors**.

Default: 1

Options: 0 to 255

Function: Indicates the priority of this neighbor, with 255 indicating the highest priority. The neighbor priority value is used in multiaccess networks for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network.

Instructions: Either accept the default neighbor priority value or enter some other value between 0 and 255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

OSPF Area Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF > Areas

Default: Enable

Options: Enable | Disable

Function: Allows you to enable and disable this area. This parameter is useful if you want to temporarily disable an area rather than delete it.

Instructions: Set this parameter to Disable if you want to disable this area. Set this parameter to Enable if you previously disabled the area and now want to re-enable it. This will cause OSPF to restart.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.2

Parameter: Authentication Type

Path: Configuration Manager > Protocols > IP > OSPF > Areas

Default: None

Options: None | Simplepassword

Function: Enables or disables password authentication for the area. If you select Simplepassword (enabling password authentication), only those routers that share the correct password will be able to communicate with each other. If you accept the default, None, password authentication is disabled for this area.

Instructions: Either accept the default value, None, to disable password authentication or select Simplepassword to enable password authentication.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.5

Parameter: Import AS Extern

Path: Configuration Manager > Protocols > IP > OSPF > Areas

Default: Yes

Options: Yes | No

Function: Indicates whether or not this area imports AS external link-state advertisements. If this area does not import AS external link-state advertisements, it is a stub area. If it does import AS external link-state advertisements, it is not a stub area.

Instructions: Set to No if this area functions as a stub area. Otherwise, accept the default value, Yes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.6

Parameter: Stub Metric

Path:	Configuration Manager > Protocols > IP > OSPF > Areas
Default:	1
Options:	1 to 255
Function:	When an area border router is connected to a stub area, it generates a default link summary into the area specifying a default route. The stub metric is the cost of that route. By default, Stub Metric equals 1. This parameter has meaning only when the Import AS Extern parameter is set to No.
Instructions:	Either accept the stub metric default value, 1, or supply the appropriate Stub Metric value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.7

Parameter: Import Summaries

Path:	Configuration Manager > Protocols > IP > OSPF > Areas
Default:	True
Options:	True | False
Function:	Specifies whether network summaries are flooded into a stub area. This variable has meaning only if the Import AS Extern parameter is set to No.
Instructions:	Set to False if Import AS Extern is set to No and you do not want network summaries imported into the stub area. Otherwise, set to True.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.8

Area Range Parameters

Parameter: Range Net

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges** > **Add**

Default: None

Options: Any network number

Function: Allows you to assign a single network address to a group of subnets. This network address, together with the subnet mask you provide, specifies the subnets to be grouped in this area range. Just one link summary advertisement will be generated for all subnets in this range, rather than one link summary advertisement for each of the subnets included in that network.

Instructions: Enter the appropriate network number in dotted decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.5

Parameter: Range Mask

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges** > **Add**

Default: None

Options: Any address mask

Function: This parameter, together with Range Net, indicates all of the networks that belong to this range. The range mask is not restricted to the natural address class mask for the address supplied at Range Net.

Instructions: Enter the appropriate subnet mask in dotted decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges**

Default: Enable

Options: Enable | Disable

Function: Enables or disables this range for the specified area. This parameter is useful if you want to disable the range, rather than delete it.

Instructions: Set this parameter to Disable if you want to disable this range. Set the parameter to Enable if you previously disabled this range and now want to re-enable it.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.2

Parameter: Mask

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges**

Default: None

Options: Any address mask

Function: This parameter allows you to change the mask portion of this area range. Mask, together with Range Net, indicates all of the networks that belong to this range. Mask is not restricted to the natural address class mask for the address supplied at Range Net.

Instructions: Enter the appropriate address mask in dotted decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Parameter: Status

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges**

Default: Advertise

Options: Advertise | Do Not Advertise

Function: Specifies whether the border router advertises a summary route to other areas.

Instructions: Select Do Not Advertise if you want to hide the existence of certain networks from other areas. By default, the border router advertises a single route for the range of routes you specify.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.7

Parameter: **Metric**

Path: Configuration Manager > Protocols > IP > OSPF > Areas > **Ranges**

Default: 0

Options: 0 to 2147483647

Function: Specifies the metric to advertise into other areas as the distance from the OSPF router to any network in the range.

Instructions: If you select 0, the router uses the value calculated by OSPF.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.8

OSPF Virtual Interface Parameters

Parameter: **Transit Area ID**

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: None

Options: Any area ID

Function: Identifies the transit area through which this virtual link is configured.

Instructions: Enter the appropriate area ID in dotted decimal notation. The transit area must contain the neighboring router identified in the Neighbors Router ID parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.6.1.4

Parameter: **Neighbor's Router ID**

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: None

Options: Any IP address

Function: Identifies the interface at the other end of this virtual link.

Instructions: Enter the appropriate IP address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.5

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: Enable

Options: Enable | Disable

Function: Enables or disables this virtual link. This parameter is useful when you want to temporarily disable a virtual link rather than delete it.

Instructions: Set to Disable to turn off this virtual link. Set to Enable if you previously disabled this virtual link and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2

Parameter: Transit Delay

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: 1 second

Options: 1 to 360 seconds

Function: Indicates the estimated number of seconds it takes to transmit a link state update packet over this interface.

Instructions: Either accept the default value of 1 second or enter a new value between 1 and 360 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.6

Parameter: Retransmit Interval

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: 5 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between link-state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets. This value should be well over the expected round-trip time. Although the default value is 5, we suggest the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 15 seconds; for point-to-multipoint, 15 seconds.

Instructions: Either accept the default value of 5 seconds or set the retransmit interval to some other value between 1 and 360 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.7

Parameter: Hello Interval

Path: Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces

Default: 15 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between the hello packets that the router sends on the interface. Although the default value is 15 seconds, we suggest the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 20 seconds, for point-to-multipoint, 15 seconds.

Instructions: Either accept the default value of 15 seconds or set the Hello Interval parameter to some other value between 1 and 360 seconds. This value must be the same for the virtual neighbor and for all routers attached to the same network.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.8

Parameter: Dead Interval

Path:	Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces
Default:	60 seconds
Options:	1 to 2000 seconds
Function:	Indicates the number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. This value should be some multiple of the Hello Interval. Although the default value is 60 seconds, we suggest the following values for this parameter: for broadcast, 40 seconds; for point-to-point, 60 seconds; for NBMA, 80 seconds; for point-to-multipoint, 60 seconds
Instructions:	Either accept the default value of 60 seconds, or enter some other value for this parameter. This value must be the same for all routers attached to the same network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

Parameter: Password

Path:	Configuration Manager > Protocols > IP > OSPF > Virtual Interfaces
Default:	None
Options:	Any ASCII text string up to eight characters long
Function:	Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. Password is valid only when Authentication Type is set to Simplepassword.
Instructions:	Enter the appropriate password. All routes in the same area must either have no authentication or have the same password.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.10

RIP Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Enable

Options: Enable │ Disable

Function: Specifies whether the Routing Information Protocol (RIP) is enabled on this interface.

Instructions: Select Enable to enable RIP on this interface. Select Disable to disable RIP on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.2

Parameter: RIP Supply

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Enable

Options: Enable │ Disable

Function: Specifies whether the interface transmits periodic RIP updates to neighboring networks.

Instructions: Select Enable to configure the interface to transmit RIP updates. Select Disable to prohibit the interface from transmitting RIP updates.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.5

Parameter: RIP Listen

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Enable

Options: Enable | Disable

Function: Specifies whether this interface listens to RIP updates from neighboring networks.

Instructions: Select Enable to configure this interface to listen to RIP updates and, thus, add received routing information to its internal routing table. If you select Enable, a configured policy can still prohibit the interface from updating its internal routing tables. Select Disable to configure the interface to ignore RIP updates from neighboring routers. Thus, the interface does not add received routing information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.6

Parameter: Default Route Supply

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Disable

Options: Enable | Disable | Generate

Function: Specifies whether or not the interface advertises a default route in RIP updates sent to neighboring networks. When a router does not know the route of a particular address, it uses the default route as the destination. A configured policy can override this setting. This parameter is independent of the RIP Supply parameter.

Instructions: If you select Enable, RIP advertises the default route if it is present in the routing table -- that is, if you have statically included a default route in the table or if the router has learned the default route (0.0.0.0) dynamically. If you select Generate, RIP advertises a default route whether or not a default route is present in the routing table. (This parameter does not cause RIP to create a routing table entry for a default route; the route will not be visible in the routing table.)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.7

Parameter: Default Route Listen

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Disable

Options: Enable | Disable

Function: Specifies whether or not IP adds default route information to its internal routing table.

Instructions: Select Enable to configure the RIP interface to listen for and potentially add the default route (0.0.0.0) information to its internal routing table. Note that you must also enable RIP Listen on this interface. A configured policy can override this setting. Select Disable to prohibit the RIP interface from adding the default route (0.0.0.0) information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.8

Parameter: Poisoned Reverse

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Poisoned

Options: Poisoned | Actual | Split

Function: Specifies how the RIP interface advertises routes it learns from an adjacent network in periodic updates subsequently sent to that network.

Instructions: Select Poisoned to configure this RIP interface to implement poisoned reverse. When poisoned reverse is enabled, the RIP interface advertises routes to the adjacent network from which it has learned the routes. In RIP updates, RIP uses a hop count of RIP Network Diameter plus one, thus declaring the destination unreachable. Poisoned reverse can speed up the convergence of the network routing tables. Select Split to configure this RIP interface to implement a split horizon. When split horizon is enabled, the RIP interface omits routes learned from a neighbor in RIP updates subsequently sent to that neighbor. Select Actual to configure this RIP interface to advertise routes with the learned cost. This is useful on a Frame Relay interface that has virtual connections (VCs) to different routers that are part of the same logical IP subnet.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.9

Parameter: Time to Live

Path	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	1
Options:	1 to 255 hops
Function:	Specifies a TTL value to be inserted in the IP header for RIP updates. Certain RIP implementations ignore packets with a TTL value of 1 hop. Use this parameter to provide interoperability with such implementations.
Instructions:	Setting a TTL of 1 prevents RIP updates from inadvertently getting off the local network. Increasing the TTL introduces the risk of the update getting off the local network and being forwarded around the network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.11

Parameter: Broadcast Timer

Path:	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	30 seconds; 1 hous for Dial Optimized Routing
Options:	5 seconds to 86,400 seconds (24 hours); 1 hour to 1,209,600 seconds (2 weeks) for Dial Optimized Routing
Function:	Specifies how frequently RIP does a full update of the routing table.
Instructions:	Enter a value in 5-second increments.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.12

Parameter: Timeout Timer

Path:	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	90 seconds; 3 hours for Dial Optimized Routing
Options:	15 seconds to 259,200 seconds (72 hours); s hours to 3,628,800 seconds (6 weeks) for Dial Optimized Routing
Function:	Specifies the time period that RIP will wait for an update for a particular network before declaring it to be unreachable.
Instructions:	We recommend a timeout value of the broadcast time multiplied by 3. Enter a time in 5-second increments.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.13

Parameter: Holddown Timer

Path:	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	90 seconds; 3 hours for Dial Optimized Routing
Options:	15 seconds to 259,200 seconds (72 hours); 3 hours to 3,628,800 seconds (6 weeks) for Dial Optimized Routing
Function:	Specifies the time period that unusable routes will be advertised through this interface after the route has become invalid.
Instructions:	We recommend a timeout value of the broadcast time multiplied by 3. Enter a time in 5-second increments. This parameter affects how long a route remains in the routing table after the route has become unusable. To guarantee the holddown time for each interface, RIP uses the largest holddown value as the amount of time to keep the route in the routing table. Please note that if a route to a destination becomes unusable, the holddown value will not affect the router's ability to learn new routes to the same destination.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.14

Parameter: RIP Mode

Path:	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	RIP I
Options:	RIP I | RIP II | RIPII with aggregation
Function:	Specifies which mode of RIP to run.
Instructions:	If you specify RIP I, RIP generates RIP Version 1 packets only. The destination IP address is the directed broadcast address, and the destination MAC address is the broadcast address. Select RIP I if any of the listening devices are RIP Version 1-only devices. If you select RIP II, RIP generates RIP Version 2 updates with the destination MAC address set to the multicast address of 224.0.0.9 specified in the RIP Version 2 RFC. The destination MAC address on Ethernet and FDDI networks will be the corresponding multicast address. On all other media, the destination MAC address will be the broadcast address. RIP does not aggregate subnet information in the updates. If you select RIPII with aggregation, RIP generates Version 2 updates but aggregates subnet information in the manner of RIP Version 1. We recommend the use of RIP II mode, with or without aggregation, rather than RIP I mode, especially if unnumbered point-to-point links or variable-length subnets are used.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.15

Parameter: Triggered Updates

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: Disable

Options: Enable │ Disable

Function: Configures RIP to generate an update each time it recalculates a route's metric.

Instructions: For compatibility with routers running Version 8.10 or earlier, disable this feature. Pre-V9.00 implementations of RIP do not support triggered updates. If you enable triggered updates, RIP will generate triggered updates with a maximum frequency of one every 5 seconds. The route will include all changes that occurred in the last 5 seconds. This enforced interval prevents RIP from monopolizing CPU resources during periods of instability.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.16

Parameter: Authentication Type

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: None

Options: None │ Simple

Function: Specifies the way RIP handles simple authentication in RIP2 mode.

Instructions: If you are running RIP in RIP2 mode and do not want authentication, set this parameter to None. If you set the parameter to Simple, RIP drops all received Version 1 updates and processes only Version 2 updates with the correct password set.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.17

Parameter: Authentication Password

Path: Configuration Manager > Protocols > IP > RIP > Interfaces

Default: None

Options: A valid password string up to 16 characters

Function: Specifies a password.

Instructions: Set the Authentication Type to Simple and enter a password.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.18

Parameter:	**Stabilization Timer**
Path:	Configuration Manager > Protocols > IP > RIP > Interfaces
Default:	120 seconds
Options:	0 to 86,400 seconds
Function:	Specifies the interval that RIP uses as its initial stabilization period.
Instructions:	Specify an interval that will allow RIP to learn all routes from its neighbors before sending a full routing update on the interface.
MIB Object ID:	

Appendix B
Routing Policies

Parameter:	**Enable**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Accept Policies
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this policy.
Instructions:	Set to Disable to disable the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.2
	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.2
	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.2
	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.2
	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.2

Parameter:	**Name**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Accept Policies
Default:	None
Options:	Any alphanumeric character string
Function:	Identifies this accept policy.
Instructions:	Specify a user name for the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.4
	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.4
	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.4
	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.4
	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.4

Parameter:	**Networks**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies
Default:	An empty list
Options:	A list of network identifiers. Each entry consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or a range of networks.
Function:	Specifies the networks to which this policy applies.
Instructions:	Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Use the default empty list to match any route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.5 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.5 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.5 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.5 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.5

Parameter:	**Action**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Accept Policies
Default:	RIP, OSPF, EGP: Accept; BGP-3, BGP-4: Ignore
Options:	Accept | Ignore
Function:	Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager.
Instructions:	Specify Accept to consider the route for insertion in the routing table. To drop the route, specify Ignore.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.6 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.6 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.6 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.6 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.6

Parameter:	**Route Preference**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies
Default:	1
Options:	1 to 16
Function:	Assigns a metric value (the higher the number, the greater the preference) to a route that the protocol forwards to the routing table manager. If confronted with multiple routes to the same destination, the routing table manager may need to use this value to decide which route to insert.
Instructions:	Either accept the default value, 1, or enter a new value. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference and routes for the most specific networks (longest address and mask) should have the highest preference.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.7 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.7 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.7 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.7 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.7

Parameter:	**Rule Precedenc**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies
Default:	0
Options:	A metric value
Function:	Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with a lower value).
Instructions:	Use this value to specify the order of precedence for policies that match the same route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.8 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.8 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.8 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.8 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.8

RIP-Specific Accept Policy Parameters

Parameter: From Gateway

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list.

Instructions: Use the default empty list to indicate that this policy applies to RIP updates from any router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.10

Parameter: Received on Interface

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP addresses of one or more interfaces on this router. This policy applies to RIP updates received on interfaces that appear on this list.

Instructions: Use the default empty list to indicate that this policy applies to RIP updates received on any interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.11

Parameter:	**Apply Subnet Mask**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Default:	Null
Options:	Null or IP address mask
Function:	Specifies a mask that will override the interface's subnet mask in the presence of networks with variable-length subnet masks.
Instructions:	Supply a mask, set the Action parameter to Accept, and use the default Network parameter (an empty list). If you specify a mask of 0.0.0.0, the router determines which mask to apply. For example, if the network in the update is a subnet of the same network as the receiving interface, the router applies the mask of the receiving interface. If the network in the update is a subnet of a different natural network, the router applies the natural mask of that network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.1.1.12

OSPF-Specific Accept Policy Parameters

Parameter:	**Type**
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Default:	Any
Options:	Type 1 | Type 2 | Any
Function:	Describes which types of OSPF ASE routes match this policy.
Instructions:	To match either Type 1 or Type 2, use the default, Any.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.3.1.10

Parameter: Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies

Default: An empty list

Options: A list of tag values

Function: Specifies OSPF tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain the tag values on this list.

Instructions: Use the default empty list to indicate that this policy applies to OSPF ASE advertisements with any tag value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.11

EGP-Specific Accept Policy Parameters

Parameter: Peer List

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP addresses of one or more EGP peers. This policy applies to EGP advertisements from the peers on this list.

Instructions: Use the default empty list to indicate that this policy applies to EGP advertisements from any EGP peer.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.10

Parameter: AS List

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous system numbers. This policy applies to EGP advertisements from peers located in the autonomous systems on this list.

Instructions: Use the default empty list to indicate that this policy applies to EGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.11

Parameter: Gateway List

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use these gateways as the next hop.

Instructions: Use the default empty list to indicate that this policy applies to EGP advertisements with any gateway address.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.12

BGP-3-Specific Accept Policy Parameters

Parameter: Injection List

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies

Default: An empty list

Options: A list of network identifiers

Function: Specifies network IDs to be included in the routing table in place of the network IDs listed in the received advertisement.

Instructions: Specify a non-null value only if the Action parameter is set to Accept. The values you enter in the injection list determine the action taken. If you supply a list of network IDs, these IDs are injected into the routing table instead of the actual received IDs.

If you use the default (an empty list), the actual received network IDs are injected into the routing table.

If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual received network IDs are injected into the routing table along with the other IDs in the injection list. This allows insertion of an aggregate or default along with the actual networks.

In the current release, the only valid network ID that you can include in an injection list is the default ID, 0.0.0.0/0.0.0.0. This parameter replaces the received routes with the default route and places the default route in the routing table. This parameter associates the default route with the attributes of the best route that matches the policy.

If you are constructing a BGP-3 or BGP-4 accept policy, keep in mind that this parameter does not perform route aggregation as defined in RFC 1654. To aggregate routes in a transit AS, you must construct an announce policy and use the announce Advertise parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.9

Parameter: Peer AS

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies

Default: An empty list

Options: A list of autonomous system numbers, each ranging from 1 to 65536

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements from peers in those ASs.

Instructions: Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.10

Parameter: Peer Address

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP peers. This policy applies to BGP advertisements from the peers on this list.

Instructions: To indicate that this policy applies to BGP advertisements from any BGP peer, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.11

Parameter: Originating AS

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements that originate from the ASs on this list.

Instructions: To indicate that the policy applies to BGP advertisements originating from any AS, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.12

Parameter: Route Origin

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Accept Policies

Default: Any

Options: Any │ IGP │ EGP │ IGP or EGP │ Incomplete │ Incomplete or IGP │ Incomplete or EGP

Function: Specifies the values of the BGP origin path attribute that apply to this policy.

Instructions: Select the origin values you wish to accept for this policy.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.13

Parameter: BGP-3 Route Preference

Path: Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies

Default: 1

Options: 1 to 16

Function: Specifies a value that is used to compare a route that matches this policy with other BGP-3 routes that match the policy. The larger the value, the greater the preference.

Instructions: To specify maximum preference, enter 16. Valid only if the Action parameter is set to Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.14

Parameter: AS Weight Class

Path: Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies

Default: Weight Class 1

Options: Weight Class 1 to Weight Class 8

Function: Indicates which weight class should be used when calculating the AS path weight.

Instructions: Set the Action parameter to Accept and supply a valid BGP-3 weight class.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.15

Parameter:	**Community Match**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Accept Policies
Default:	An empty list
Options:	A list of BGP communities
Function:	Specifies one or more BGP communities. This policy applies to all BGP advertisements that match the list.
Instructions:	Supply an octet string using the following format: each community ID is 4 bytes long; 0 in the two most significant bits causes the router to perform the match on the lower 16 bits; the default empty list means “match any list.”
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.15

BGP-4-Specific Accept Policy Parameters

Parameter: Injection List

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of network identifiers

Function: Specifies network IDs to be included in the routing table in place of the network IDs listed in the received advertisement.

Instructions: Specify a non-null value only if the Action parameter is set to Accept. The values you enter in the injection list determine the action taken.

If you supply a list of network IDs, these IDs are injected into the routing table instead of the actual received IDs.

If you use the default (an empty list), the actual received network IDs are injected into the routing table.If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual received network IDs are injected into the routing table along with the other IDs in the injection list. This allows insertion of an aggregate or default along with the actual network.

In the current release, the only valid network ID that you can include in an injection list is the default ID, 0.0.0.0/0.0.0.0. This parameter replaces the received routes with the default route and places the default route in the routing table.This parameter associates the default route with the attributes of the best route that matches the policy.

If you are constructing a BGP-3 or BGP-4 accept policy, keep in mind that the Injection List parameter does not perform route aggregation as defined in RFC 1654. To aggregate routes in a transit AS, you must construct an announce policy and use the announce Advertise parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.9

Parameter: Peer AS

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of autonomous system numbers, each ranging from 1 to 65536

Function: Specifies one or more ASs. This policy applies to BGP advertisements from peers in the autonomous systems on this list.

Instructions: Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.10

Parameter: Peer Address

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP peers. This policy applies to BGP advertisements from the peers on this list.

Instructions: To indicate that this policy applies to BGP advertisements from any BGP peer, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.11

Parameter: Originating AS

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements that originate from the ASs on this list.

Instructions: To indicate that the policy applies to BGP advertisements originating from any AS, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.12

Parameter: Route Origin

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: Any

Options: Any | IGP | EGP | IGP or EGP | Incomplete | Incomplete or IGP | Incomplete or EGP

Function: Specifies which values of the BGP origin attribute apply to this policy.

Instructions: Select the origin values you wish to accept for this policy.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.13

Parameter: Aggregator AS List

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of AS numbers

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements that contain in their Aggregator path attribute an AS number on this list.

Instructions: To specify that the policy applies to BGP advertisements with any AS number in the Aggregator path attribute, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.14

Parameter: Aggregator Router List

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP routers. This policy applies to BGP advertisements that contain in their Aggregator path attribute an IP address on this list.

Instructions: To specify that this policy applies to BGP advertisements with any router address in the Aggregator path attribute, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.15

Parameter: Local Preference

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: 0

Options: 0 to 4294967295

Function: Assigns a local preference value to a route matching this policy. This value overrides the calculated value for EBGP routes or the Local Preference path attribute for IBGP routes.

Instructions: To indicate a preference, enter a value from 1 to 4294967295.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.16

Parameter: BGP-4 Preference

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: 1

Options: 1 to 16

Function: Specifies a value that can be used to compare a route that matches this policy with other BGP-4 routes. The larger the value, the greater the preference.

Instructions: To indicate maximum preference, enter 16. This parameter is valid only if the Action parameter is set for Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.17

Parameter: AS Weight Class

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies

Default: Weight Class 1

Options: Weight Class 1 to Weight Class 8

Function: Indicates which weight class value should be used when calculating the AS path weight.

Instructions: Enter a valid BGP-4 weight class. Valid only if the Action parameter is set for Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.18

Parameter: AS Pattern

Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies
Default:	Empty string
Options:	Any regular expression or empty string
Function:	Allows AS_PATH pattern matching
Instructions:	Enter a valid regular expression to indicate an AS and its position in a path. The policy applies to all routes whose AS path includes the AS in that position. For example, the expression * 200 $ means that the policy applies to all routes whose AS_PATH attribute contains AS 200 as the last AS in the path.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.9.1.19

Parameter: Community Match

Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Accept Policies
Default:	An empty list
Options:	A list of BGP communities
Function:	Specifies one or more BGP communities. This policy applies to all BGP advertisements that match the list.
Instructions:	Supply an octet string using the following format: each community ID is 4 bytes long; 0 in the two most significant bits causes the router to perform the match on the lower 16 bits; the default empty list means “match any list.”
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.9.1.20

IP Announce Policy Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this policy.
Instructions:	Set to Disable to disable the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.2
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.2
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.2
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.2
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.2

Parameter:	**Name**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	None
Options:	Any alphanumeric character string
Function:	Identifies this policy.
Instructions:	Enter a unique name for the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.4
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.4
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.4
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.4
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.4

Parameter:	**Networks**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of network identifiers. Each identifier consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or a range of networks
Function:	Specifies which networks will match this policy.
Instructions:	Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Enter an empty list to match any route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.5
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.5
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.5
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.5
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.5

Parameter:	**Action**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	RIP, OSPF, EGP: Propagate; BGP-3, BGP-4: Ignore
Options:	Propagate | Ignore
Function:	Specifies whether or not to advertise a route that matches this policy.
Instructions:	To advertise the route, specify Propagate. To drop the route, specify Ignore.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.6
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.6
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.6
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.6
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.6

Parameter:	**Rule Precedence**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	0
Options:	A metric value
Function:	Specifies a metric value to be used to compare this policy with other policies that a route may match. A policy with a higher metric takes precedence over a policy with a lower metric. In case of a tie, the protocol uses an internal index value assigned to the policy by IP software. (In general, the index value is indicated by the position of the policy in the Site Manager display -- the last policy in the display has the highest index value.)
Instructions:	Use this parameter to assign precedence to policies that match the same route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.7
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.7
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.7
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.7
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.7

Parameter: Route Source

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Any

Options: Any | Direct | Static | RIP | OSPF (not valid for OSPF) | EGP | BGP

Function: Specifies one or more route source identifiers. If you select a route source ID, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: To specify any source, use the default.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.8

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.8

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.8

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.8

MIB Object ID: BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.8

Parameter:	**Advertise**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of network identifiers
Function:	Specifies network IDs to include in place of the network IDs listed in the route to be advertised.
Instructions:	Specify a non-null value only if the announce Action parameter is Propagate. The values you enter in the advertise list determine the action taken. If you supply a list of network IDs, these IDs are advertised instead of the actual IDs in the route. If you use the default (an empty list), the actual IDs are advertised. Note that by default, BGP-4 aggregates subnets into their natural network IDs. If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual network IDs are advertised along with the other IDs in the advertise list.This allows advertisement of an aggregate or default along with the actual network. If the actual network is a subnet (and the advertising protocol supports subnet advertisements), the subnet is advertised.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.10
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.10
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.10
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.10
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.10

Parameter:	**From RIP Gateway**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list. Applicable only for RIP-sourced routes and if RIP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates from any router.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.11
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.11
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.11
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.11
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.11

Parameter:	**Received on RIP Interface**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the addresses of one or more interfaces on this router. This policy applies to RIP advertisements received on the interfaces in this list. Applicable only for RIP-sourced routes and if RIP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates received on any interface.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.12
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.12
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.12
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.12
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.12

Parameter: From OSPF Router ID

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IDs of one or more OSPF routers. This policy applies to OSPF advertisements authored by a router on this list. Applicable only for OSPF-sourced routes and if OSPF is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to OSPF updates from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.13

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.13

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.13

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.13

MIB Object ID: BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.13

Parameter:	**Received OSPF Type**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies
Default:	Any
Options:	Type 1 | Type 2 | External | Internal | Any
Function:	Specifies which types of OSPF routes match this policy. Applicable only for OSPF-sourced routes and if OSPF is included as a route source.
Instructions:	To match any route type, enter Any. To match any non-ASE route, enter Internal. To match any ASE route, enter External. To match any external type 1 route, enter Type 1. To match any external type 2 route, enter Type 2.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.14
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.14
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.14
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.14
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.14

Parameter: Received OSPF Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of tag values

Function: Specifies tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain tag values in this list. Applicable only for OSPF-sourced ASE routes and if OSPF is included as a route source.

Instructions: Specify one or more tag values. Use the default empty list to indicate that this policy applies to OSPF ASEs with any tag value.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.15

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.15

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.15

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.15

MIB Object ID: BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.15

Parameter:	**From EGP Peer**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP address of one or more EGP peers. This policy applies to EGP advertisements authored by a router on this list. Applicable only for EGP source routes and if EGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements from any router.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.16
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.16
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.16
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.16
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.16

Parameter: From EGP AS

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous system numbers. This policy applies to EGP advertisements received from EGP peers in an AS on this list. Applicable only for EGP-sourced routes and if EGP is included as a route source.

Instructions: Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to EGP advertisements from peers in any AS.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.17

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.17

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.17

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.17

MIB Object ID: BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.17

Parameter:	**Received EGP Gateway**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use a gateway on this list as the next hop. Applicable only for EGP-sourced routes and if EGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements with any gateway address.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.18
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.18
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.18
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.18
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.18

Parameter:	**From BGP Peer**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP address of one or more BGP peers. This policy applies to BGP advertisements authored by a router on this list. Applicable only for BGP-sourced routes and if BGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements from any router.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.19
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.19
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.19
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.19
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.19

Parameter:	**From BGP AS**
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocol**s** > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of autonomous system numbers
Function:	Specifies one or more autonomous system numbers. This policy applies to BGP advertisements received from BGP peers in an AS on this list. Applicable only for BGP-sourced routes and if BGP is included as a route source.
Instructions:	Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.20
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.20
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.20
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.20
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.20

Parameter:	**Received BGP Next Hop**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies one or more IP addresses. This policy applies to BGP advertisements whose Next Hop attribute matches an IP address on this list. Applicable only for BGP-sourced routes and if BGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements with any Next Hop attribute.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.21
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.21
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.21
MIB Object ID:	BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.21
MIB Object ID:	BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.21

RIP-Specific Announce Policy Parameters

Parameter: External Route Source

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Default: Any

Options: Direct | Static | RIP | OSPF (with Type 2 metric) | EGP | BGP | Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.9

Parameter: Outbound Interfaces

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies a list of outbound RIP interfaces. If an interface appears in this list, the policy applies to RIP advertisements sent via that interface.

Instructions: Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to any outbound RIP interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.22

Parameter: RIP Metric

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Default: 0

Options: 0 or an export metric

Function: Specifies an optional export RIP metric to use when advertising a route that matches this policy.

Instructions: Set the Action parameter to Announce. If you use the default, the RIP metric is the routing table metric calculated for RIP plus the interface cost.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.23

OSPF-Specific Announce Policy Parameters

Parameter: Type

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: 0

Options: Type 1 | Type 2 | 0

Function: Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy.

Instructions: Enter 0 if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1. Set the Action parameter to propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.22.

Parameter: Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: Null

Options: Null or a tag value

Function: Specifies a value for the OSPF external route tag field. If the outgoing route matches this policy, the router places this value in the field.

Instructions: Set the Action parameter to Propagate and set the Automatic Tag parameter to Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.23

Parameter: Automatic Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: Disable

Options: Enable │ Disable

Function: Enables BGP/OSPF automatic tag generation.

Instructions: Select Disable (the default) to use the value you specify with the Tag parameter. Select Enable to generate a tag according to the criteria in RFC 1403 (or any superseding RFC). This parameter overrides the Tag Generation Method parameter on the OSPF Global Parameters window.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.24

Parameter: OSPF Metric

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: 0

Options: 0 or an export metric

Function: Specifies an optional OSPF metric to use when advertising a route that matches this policy.

Instructions: Set the Action parameter for Announce. If you use the default, the OSPF metric is the routing table metric.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.25

EGP-Specific Announce Policy Parameters

Parameter: External Route Source

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Default: Any

Options: Direct | Static | RIP | OSPF (with Type 2 metric) | EGP | BGP | Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.9

Parameter: EGP Peer List

Path: Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies a list of IP addresses of EGP peers. If a peer appears in this list, the policy applies to EGP advertisements sent to that peer.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that the policy applies to any BGP peer.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.22

Parameter: EGP Interface List

Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies a list of outgoing interfaces. If an interface appears on this list, the policy applies to EGP advertisements sent via that interface.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to any outbound interface.
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.23

Parameter: EGP Metric

Path:	Configuration Manager > Protocols > IP > Policy Filters > EGP > Announce Policies
Default:	0
Options:	0 or an export metric value
Function:	Specifies an optional export metric to use when advertising a route that matches this policy.
Instructions:	Select the default to indicate that the routing table metric calculated for EGP is to be used. This parameter is valid only if the Action parameter is set to Propagate.
MIB Object ID:	EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.24

BGP-3-Specific Announce Policy Parameters

Parameter: External Route Source

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: Any

Options: Direct │ Static │ RIP │ OSPF (with Type 2 metric) │ EGP │ BGP │ Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF external routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.9

Parameter: Outbound Peer AS List

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: An empty list

Options: A list of AS numbers

Function: Specifies a list of autonomous system numbers. If an AS number is included in this list, this policy applies to BGP advertisements being sent to BGP peers in that AS.

Instructions: Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to BGP advertisements going to peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.22

Parameter: Outbound Peers

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: An empty list

Options: A list of IP numbers

Function: Specifies the IP address of one or more BGP peers. If a BGP peer is included in this list, this policy applies to BGP advertisements being sent to that peer.

Instructions: Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to BGP advertisements being sent to any peer.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.23

Parameter: Inter-AS Metric Selector

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: None

Options: None | Specified | Originating

Function: Indicates whether or not an inter-AS metric is to be advertised for a network matching this policy and, if advertised, what value to use.

Instructions: Select None to indicate that no metric is to be advertised. Select Specified to indicate that the value you specify in the Specific Inter-AS Metric parameter is to be used. Select Originating to indicate that the metric from the originating protocol will be used. This parameter is valid only if the Action parameter is set to Propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.24

Parameter: **Specific Inter-AS Metric**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: Null

Options: Null or an AS metric

Function: Specifies a value for the inter-AS metric.

Instructions: Supply a value and set the inter-AS Metric Selector parameter to Specified.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.25

Parameter: **Origin**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: As Is

Options: As Is │ IGP │ EGP │ Incomplete

Function: Specifies an Origin attribute override. The Origin attribute of a route matching this policy will be replaced with the indicated value.

Instructions: To allow the existing Origin attribute, use the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.26

Parameter: **AS Path Override**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies

Default: An empty list

Options: A list of AS numbers

Function: Specifies an AS path override.

Instructions: Enter a non-null value to override the AS path attribute of a route matching this policy. Each element of the AS path is an AS number. Valid only if the Action parameter is set to Propagate. Use the default empty list to allow the existing AS path attribute to remain in the route.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.27

Parameter: Next Hop

Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Default:	Null
Options:	An IP address
Function:	Overrides the Next Hop path attribute with the IP address you specify.
Instructions:	To allow the existing Next Hop attribute, use the default null value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.8.1.28

Parameter: Community Match

Path:	Configuration Manager > Protocols > IP > Policy Filters > BGP3 > Announce Policies
Default:	An empty list
Options:	A list of BGP communities
Function:	Specifies one or more BGP communities. This policy applies to all BGP advertisements that match the list.
Instructions:	Supply an octet string using the following format: each community ID is 4 bytes long; 0 in the two most significant bits causes the router to perform the match on the lower 16 bits; the default empty list means “match any list.”
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.8.1.29

BGP-4-Specific Announce Policy Parameters

Parameter: External Route Source

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Any

Options: Direct | Static | RIP | OSPF (with Type 2 metric) | EGP | BGP | Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.9

Parameter: Outbound Peer AS

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of AS numbers

Function: Specifies a list of autonomous system numbers. If an AS number is included in this list, this policy applies to BGP advertisements being sent to BGP peers in that AS.

Instructions: Specify one or more AS numbers. Configure an empty list to indicate that this policy applies to BGP advertisements going to peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.22

Parameter: Outbound Peers

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more BGP peers. If a BGP peer is included in this list, this policy applies to BGP advertisements being sent to that peer.

Instructions: Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to BGP advertisements being sent to any peer.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.23

Parameter: Multi-Exit Discriminator

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: None

Options: None | Specified | Originating

Function: Indicates whether or not a Multi-Exit Discriminator metric is to be advertised for a network matching this policy and, if advertised, what value to use.

Instructions: Select None to indicate that no value is to be advertised. Select Specified to indicate that the value you specify for the Multi-Exit Discriminator Value parameter is to be used. Select Originating to indicate that the metric from the originating protocol is to be used. This parameter is valid only if the Action parameter is set for Propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.24

Parameter: Multi-Exit Discriminator Value

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Null

Options: Null or a metric value

Function: Specifies a metric for the Multi-Exit Discriminator attribute.

Instructions: To advertise a multi-exit discriminator value, set the Action parameter to Propagate and set the Multi-Exit Discriminator parameter to Specified.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.25

Parameter: Origin

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: As Is

Options: As Is │ IGP │ EGP │ Incomplete

Function: Specifies an Origin attribute override. The Origin attribute of a route matching this policy will be replaced with the indicated value.

Instructions: To allow the existing Origin attribute, use the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.26

Parameter: AS Path

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Null

Options: An AS path

Function: Specifies an AS path that overrides the AS-path attribute of a route matching this policy.

Instructions: Constructs a BGP-4 AS path composed of AS path segments. Each AS path segment includes a path segment type, a path segment length specifying the number of ASs in the segment, and a path segment value containing one or more AS numbers. There are two AS path segment types: Type 1, an unordered set of ASs that a route in the UPDATE message has traversed, and Type 2, an ordered set of ASs that a route in the UPDATE message has traversed.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.27

Parameter: Local Preference Override

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Local Pref Override False

Options: False | True

Function: Indicates whether or not you are supplying an override value for the Local Preference path attribute in the routing Update message. (The Local Pref attribute is valid only in an Update advertised to an IBGP peer.) If you select False, the router uses the IP route weight value to calculate the LOCAL_PREF path attribute.

Instructions: To override the Local Preference attribute, select True and supply a value for the Local Preference Value parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.28

Parameter: **Local Preference Value**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Null

Options: Null or a route weight value

Function: Specifies an override value for the Local Preference attribute.

Instructions: Enter a value and set the Local Preference Override parameter to True.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.29

Parameter: **Next Hop**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Null

Options: An IP address

Function: Overrides the Next Hop path attribute with the IP address you specify.

Instructions: To allow the existing Next Hop attribute, use the default null value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.30

Parameter: **Atomic**

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce PoliciesConfiguration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Automatic

Options: Automatic │ Force │ Ignore

Function: Allows control over the Atomic path attribute.

Instructions: By default, the router automatically sets this parameter if it knows that certain networks in aggregate range have not been included in an aggregate advertisement.

MIB Object ID: To include the Atomic attribute even if the router does not assume one is required, set the parameter to Force.

Parameter: AS Pattern

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: Empty string

Options: Any regular expression or empty string

Function: Allows AS_PATH pattern matching.

Instructions: Enter a valid regular expression to indicate an AS and its position in a path. The policy applies to all routes whose AS path includes the AS in that position. For example, the expression * 200 $ means that the policy applies to all routes whose AS_PATH attribute contains AS 200 as the last AS in the path.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.32

Parameter: Community Match

Path: Configuration Manager > Protocols > IP > Policy Filters > BGP4 > Announce Policies

Default: An empty list

Options: A list of BGP communities

Function: Specifies one or more BGP communities. This policy applies to all BGP advertisements that match the list.

Instructions: Supply an octet string using the following format: each community ID is 4 bytes long; 0 in the two most significant bits causes the router to perform the match on the lower 16 bits; the default empty list means “match any list.”

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.33

Appendix C
Import and Export Route Filters

RIP

RIP Import Filters

Parameter:	**Import Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If this field is set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate network address in dotted decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.3

Parameter: **Import Mask**

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses this filter acts upon.

Instructions: For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the Net ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the Net ID and Subnet ID portions of the address will be filtered. If the Import Address field is set to 0.0.0.0, and this parameter is set to 0.0.0.0, then the filter applies to *all* routes. If the Import Address field is set to 0.0.0.0, and this parameter is set to 255.255.255.255, then the filter applies to the *default* route. Enter the mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.4

Parameter: **RIP Gateway**

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 0.0.0.0

Options: Any IP address

Function: Identifies, by IP address, the router that is sending the updates. This filter will apply to updates from that router. If this field is set to 0.0.0.0, the filter applies to updates from any router.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.7

Parameter: Interface

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default: 0.0.0.0
Options: Any IP address
Function: Specifies the local IP address of the interface that connects this router to the RIP gateway. This filter will apply only to those updates received on this interface. If set to 0.0.0.0, this filter applies to all interfaces.
Instructions: Enter the appropriate IP address in dotted decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.8

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default: Accept
Options: Accept | Ignore
Function: Specifies whether the route is transferred to the routing tables. If this parameter is set to Accept (default), the routing information is sent to the routing tables. If this parameter is set to Ignore, the routing information is dropped.
Instructions: Either accept the default Accept, or select Ignore.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.5

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 1

Options: 1 to 16

Function: Assigns a weighted preference value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.

Instructions: If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this RIP-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference). Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference. The default preference for static routes is 1, but may be set to any value between 1 and 16 (refer to "Editing Static Route Parameters" for more information). If you want to grant a RIP-derived route preference over a static route, make sure the preference value you assign to the RIP-derived route is greater than the preference value of the static route you want it to override.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.6

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: Enable

Options: Enable | Disable

Function: Enables or disables this import route filter.

Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.2

Parameter:	**Apply Subnet Mask**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	Null or IP address mask
Options:	Specifies a mask that will override the interface's subnet mask in the presence of networks with variable-length subnet masks.
Function:	Supply a mask, set the Action parameter to Accept, and use the default Network parameter (an empty list).
Instructions:	If you specify a mask of 0.0.0.0, the router determines which mask to apply. For example, if the network in the update is a subnet of the same network as the receiving interface, the router applies the mask of the receiving interface. If the network in the update is a subnet of a different natural network, the router applies the natural mask of that network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.9

RIP Export Filters

Parameter:	**Export Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate IP address in dotted decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.9.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts

Instructions: For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the subnet ID, and the final 8 bits to the host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the Net ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the net ID and subnet ID portions of the address will be filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.4

Parameter: From Protocol

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters

Default: Any

Options: Any | RIP | EGP | OSPF | Direct | Static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP-, OSPF-, EGP-, or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.5

Parameter: Interface

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: 0.0.0.0
Options: Any IP address
Function: Identifies the outgoing IP interface for the RIP update. This filter will only apply to this interface. If set to 0.0.0.0, this filter applies to all interfaces.
Instructions: Enter the appropriate IP address in dotted decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.7

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: Enable
Options: Enable | Disable
Function: Enables or disables this export route filter.
Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.2

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: Propagate
Options: Propagate | Ignore | Aggregate
Function: Controls the flow of routing information. If you set this parameter to Propagate, this route is advertised. If you set this parameter to Ignore, advertising of this route is suppressed. If you set this parameter to Aggregate, the network is not explicitly advertised. Instead, the default route (0.0.0.0) is advertised.
Instructions: Either accept the default, Propagate, or select Ignore or Aggregate.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.6

Parameter:	**RIP Metric**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default:	0 (the actual route cost as learned)
Options:	0 to 15
Function:	Assigns a RIP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used.
Instructions:	Accept the default value 0 or enter a new value. Do not use a value that exceeds the diameter of the RIP network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.9.1.8

OSPF Route Filters

OSPF Import Filters

Parameter:	**Import Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Import Filters
Default:	None
Options:	An IP address
Function:	Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate network address in dotted decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.10.1.3

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Import Filters

Default: 0

Options: 0 to 16

Function: Assigns a weighted preference value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.

Instructions: If this hierarchy is acceptable, accept the default value 0 for preference. If you want to grant preference to this OSPF-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference). Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference. The default preference for static routes is 0, but it may be set to any value between 0 and 16 If you want to grant an OSPF-derived route preference over a static route, make sure the preference value you assign to the OSPF-derived route is greater than the preference value of the static route you want it to override.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.6

OSPF Export Filters

Parameter: Export Address

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: 0.0.0.0

Options: Any IP network address

Function: Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: None

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts.

Instructions: For example, consider Class B Network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the Net ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the Net ID and Subnet ID portions of the address will be filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.4

Parameter: Export From Protocol

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: RIP

Options: Any | RIP | EGP | OSPF | Direct | static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP, EGP, OSPF, or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.5

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: Enable

Options: Enable | Disable

Function: Enables or disables this export route filter.

Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.2

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: Propagate

Options: Propagate | Ignore

Function: Controls the flow of routing information. If you set this parameter to Propagate, this route is advertised. If you set this parameter to Ignore, advertising of this route is suppressed.

Instructions: Either accept the default Propagate, or select Ignore.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.6

Parameter: Type

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: Type 1

Options: As Is, Type 1, Type 2

Function: Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy.

Instructions: Select As Is if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1. Set the Action parameter to Propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.7

Parameter: Tag

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: 1

Options: 1 to 2147483647

Function: Sets the tag value for the AS external advertisement that is generated for this network. This parameter has meaning only when the Action parameter is set to Propagate.

Instructions: Enter the appropriate tag.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.8

Parameter: Auto Tag

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: Disable

Options: Enable | Disable

Function: If enabled, the router creates a tag for this route as described in RFC 1364 (BGP/OSPF Interaction).

Instructions: Set to Enable if you are running BGP-3 as your exterior gateway protocol.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.9

BGP-3

BGP-3 Import Filters

Parameter: Import Address

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 0.0.0.0

Options: Any IP network address

Function: Identifies, by IP address, the network to which this filter applies.

Instructions: Enter the appropriate network address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.3

Parameter: Import Mask

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts.

Instructions: For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the Net ID portion of the address will be filtered. If you enter the mask 255.255.255.0, the net ID and subnet ID portions of the address will be filtered. If you set the Import Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Import Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.4

Parameter: Import Peer AS

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 0

Options: 0 to 65535

Function: Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs. This filter will apply to updates from this router. The value 0 means “any” AS.

Instructions: Enter the appropriate AS number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.7

Parameter: Import Peer Address

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 0.0.0.0

Options: Any IP address

Function: Specifies the IP address of the interface on the remote side of this BGP peer connection. This filter will apply to updates from this router. The value 0 means “any” peer.

Instructions: Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.8

Parameter: Import Originating AS

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 0

Options: 0 to 65535

Function: Specifies the AS from which the route originated (the last AS in the AS path). The filter will apply to updates created by any routers in this AS. The value 0 means “any” AS.

Instructions: Enter the appropriate AS number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.9

Parameter: Import Route Origin

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: Any

Options: Any | IGP | EGP | Incomplete

Function: Specifies the value of the Origin Path attribute in the Update message received.

Instructions: Set the appropriate Import Route Origin value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.10

Parameter: Import Action

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: Ignore

Options: Accept | Ignore

Function: Specifies whether the route is transferred to the routing tables. If you set this parameter to Accept, the routing information is sent to the routing tables. If you select Ignore, the routing information is dropped.

Instructions: Either accept the default, Ignore, or select Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.5

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: Enable

Options: Enable | Disable

Function: Enables or disable this import route filter.

Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.2

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 1

Options: 1 to 16

Function: Assigns a weighted preference value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, BGP-3, static, OSPF, external, and RIP. If Intra-AS IBGP routing is used, then any other route source is preferred over a BGP-3 route.

Instructions: If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this BGP-3-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference). The default preference for static routes is 16, but may be set to any value between 1 and 16. If you want to grant a BGP-3-derived route preference over a static route, make sure the preference you assign to the BGP-3-derived route exceeds the preference value of the static route you want it to override. Either accept the default value, 1, or enter a new value. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.11

Parameter: BGP-3 Preference

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Import Filters

Default: 1

Options: 1 to 2147483647

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple BGP-3 routes to the same destination, the router, by default, grants preference to routes assigned the highest preference value.

Instructions: Either accept the default value, 1, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.12

BGP-3 Export Filters

Parameter: Export Address

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters

Default: 0.0.0.0

Options: Any IP network address

Function: Identifies, by IP address, the network to which this filter applies. If this field is left blank, the filter applies to all networks.

Instructions: Enter the appropriate network address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts.

Instructions: For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the Net ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the Net ID and Subnet ID portions of the address will be filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.4

Parameter:	**Export from Protocol**
Path:	Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default:	Any
Options:	Any | RIP | EGP | OSPF | Direct | Static | BGP-3
Function:	Identifies the source of the routing information: direct connection, static route, or a RIP-, EGP-, OSPF-, or BGP-3 derived route.
Instructions:	Select the appropriate option.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.15.1.5

Parameter:	**Export Peer AS**
Path:	Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default:	0
Options:	1 to 65535
Function:	Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs. This filter will apply to updates sent to any router in this AS. The value 0 means “any” AS.
Instructions:	Enter the appropriate AS number.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.15.1.7

Parameter:	**Export Peer Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default:	0.0.0.0
Options:	Any IP address
Function:	Specifies the IP address of the interface on the remote side of this BGP peer connection. This filter will apply to updates sent to this router. The value 0.0.0.0 means “any” peer.
Instructions:	Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.15.1.8

Parameter: Export Enable

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default: Enable
Options: Enable | Disable
Function: Enables or disables this export route filter.
Instructions: Set to Disable if you want to disable this filter. Set to Enable if you want to enable this filter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.2

Parameter: Export Action

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default: Ignore
Options: Propagate | Ignore | Aggregate
Function: Controls the flow of routing information. If set to Propagate, this route is advertised. If set to Ignore, advertising of this route is suppressed. If set to Aggregate, the network is not explicitly advertised. Instead, the default route (0.0.0.0) is advertised.
Instructions: Select Propagate, Ignore, or Aggregate.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.6

Parameter: Export Use Inter AS Metric

Path: Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default: None
Options: None | Specified | Originating
Function: Specifies whether or not an Inter AS metric is advertised for the associated networks. If set to None, then no metric is advertised. If set to Specified, then the value specified for the Export Inter AS Metric parameter is advertised. If set to Originating, then the metric from the originating protocol is advertised. This parameter is only valid if Export Action is set to propagate.
Instructions: Set to the appropriate option.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.11

Parameter:	**Export Origin**
Path:	Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default:	Any
Options:	Any | IGP | EGP | Incomplete
Function:	If From Protocol is set to RIP or Static, and Action is set to Propagate, you can use this parameter to change the Origin attribute that is advertised for this network.
Instructions:	If you want to change the Origin attribute, select a valid option.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.15.1.13

Parameter:	**Export Neighbor AS**
Path:	Configuration Manager > Protocols > IP > Route Filters > BGP3 > Export Filters
Default:	0
Options:	0 to 65535
Function:	If the Export Action parameter is set to Propagate, and the Export Origin parameter is set to EGP, then this parameter must be set to a nonzero value. The value specified here is used as the EGP neighbor AS number when the AS path is constructed.
Instructions:	Specify a value within the assigned range.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.15.1.14

EGP

EGP Import Filters

Parameter:	**Import Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If this field is set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate network address in dotted decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.12.1.3

Parameter:	**Import Peer**
Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters
Default:	0.0.0.0
Options:	Any IP address
Function:	Specifies the IP address of the interface on the remote side of this EGP peer connection. This filter will apply to updates from this router. The default 0.0.0.0 means “any” peer.
Instructions:	Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.12.1.7

Parameter: Import Autonomous System

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters

Default: 0

Options: 0 to 65536

Function: Identifies the AS to which the EGP router at the remote end of this EGP peer connection belongs. This filter will apply to updates from this router. The default 0 means “any” AS.

Instructions: Enter the appropriate AS number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.12.1.8

Parameter: Import Gateway

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters

Default: 0.0.0.0

Options: Any IP address

Function: Specifies the gateway advertised as the next hop for the network. The default value of 0 means “any” gateway.

Instructions: Enter the appropriate gateway number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.12.1.9

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters

Default: Enable

Options: Enable | Disable

Function: Enables or disables this import route filter.

Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.12.1.2

Parameter:	**Action**
Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters
Default:	Accept
Options:	Accept | Ignore
Function:	Specifies whether the route is transferred to the routing tables. If you select Accept (default), the routing information is sent to the routing tables. If you select Ignore, the routing information is dropped.
Instructions:	Either accept the default, Accept, or select Ignore.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.12.1.5

Parameter:	**Preference**
Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Import Filters
Default:	1
Options:	1 to 15
Function:	Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, and RIP. If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this OSPF-derived route, assign a new preference value in the range of 1 to 15 (the greater the number, the higher the preference).
Instructions:	Either accept the default value 1, or enter a new value. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.12.1.6

EGP Export Filters

Parameter: Export Address

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters

Default: 0.0.0.0

Options: Any IP network address

Function: Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses this filter acts upon. For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. Thus, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the net ID portion of the address is filtered. If you enter the mask 255.255.255.0, the Net ID and Subnet ID portions of the address are filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.4

Parameter: Export From Protocol

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters

Default: Any

Options: Any | RIP | EGP | OSPF | Direct | Static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP-, EGP-, OSPF-, or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.5

Parameter: Export Peer

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters

Default: 0.0.0.0

Options: Any IP address

Function: Specifies the IP address of the interface on the remote side of this EGP peer connection. This filter will apply to updates from this router. The default value 0.0.0.0 means “any” router.

Instructions: Enter the IP address in dotted decimal notation. The address must be on the same subnet as a local interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.7

Parameter: Export OSPF Type

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters

Default: None

Options: Type 1 | Type 2 | Internal

Function: Specifies the type of routes to which this filter applies. If you specify Type 1, then only AS External Type 1 routes are filtered. If you specify Type 2, then only AS External Type 2 routes are filtered. Note that this parameter is only used if the Export From Protocol parameter is set to OSPF.

Instructions: Depending on the type of routes you want to filter, select Type 1, Type 2, or Internal.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.8

Parameter: Export OSPF Tag

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters
Default: 0
Options: 0 to 2147483647
Function: Specifies the tag with which this route filter is concerned. Each AS External Advertisement contains a Tag field. If the Tag field matches Import Tag, the appropriate action is taken; either the route is accepted or ignored. Note that this parameter is only used if the Export From Protocol parameter is set to OSPF.
Instructions: Enter the appropriate tag number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.9

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters
Default: Enable
Options: Enable | Disable
Function: Enables or disables this export route filter.
Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.2

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters
Default: Propagate
Options: Propagate | Ignore
Function: Controls the flow of routing information. If you select Propagate, this route is advertised. If you select Ignore, advertising of this route is suppressed.
Instructions: Either accept the default, Propagate, or select Ignore.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.6

Parameter: Interface

Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters
Default:	0.0.0.0
Options:	Any IP address
Function:	Specifies the outbound interface on which to apply this filter.
Instructions:	Specify the IP address of the interface on which you want to apply this filter.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.13.1.10

Parameter: Metric

Path:	Configuration Manager > Protocols > IP > Route Filters > EGP > Export Filters
Default:	0 (the actual route cost as learned)
Options:	0 to 255
Function:	Assigns an EGP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used.
Instructions:	Either accept the default metric value, 0, or enter a new value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.13.1.11

Appendix D
Route Weight Worksheet

1. **Select one route from the following list:**

Direct	0
OSPF Internal	0
OSPF External	16 (OSPF Import Preference)
RIP	16 (RIP Import Preference)
EGP	(16 - EGP Import Preference)
BGP	(16 - BGP Import Preference)
Static	(16 - SR Preference)

2. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   134217728 x _____=_____

   or

   0x8000000 x ____ = _____

3. **Select one route from the following list:**

Direct	0
OSPF Intra-Area Internal	0
OSPF Inter-Area Internal	1
OSPF Type 1 External	2
OSPF Type 2 External	3
Non-OSPF External (RIP, EGP, BGP)	3
Static	3

4. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   16777216 x _____ = _____

   or

   0x1000000 x _____ = _____

5. **Select one route from the following list:**

Direct	0
OSPF Internal	0
OSPF Type 2 External w/ASE Metric support enabled	0
EBGP	2
RIP	4
EGP	5
Static	6
OSPF Type 2 External w/ASE Metric support disabled	7

6. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   2097152 x _____ = _____

   or

   0x200000 x _____ = _____

7. **Select one route from the following list and calculate the associated value using the formulas supplied:**

Direct	OSPF Interface ? OSPF Interface metric : 0
OSPF imported	OSPF-TOS-0 Metric
BGP-3 imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP3 Import BGP3 Preference) + (IGP Origin ? 0 : 4096) + (AS Weighted Path Length <= 4095 ? AS Weighted Path Length : 4095) or (0x2000 * (16 - BGP3 Import BGP3 Preference) + (IGP Origin ? 0 : 0x1000) + (AS Weighted Path Length <= 0x0fff ? AS Weighted Path Length : 0x0fff)
BGP-4 imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP4 Import BGP3 Preference)) + (IGP Origin ? 0 : 4096) + (AS Weighted Path Length <= 4095 ? AS Weighted Path Length : 4095) or 0x2000 * (16 - BGP4 Import BGP3 Preference) + (IGP Origin ? 0 : 0x1000) + (AS Weighted Path Length <= 0x0fff ? AS Weighted Path Length : 0x0fff)

IBGP imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP4 Import BGP3 Preference) + (Local Pref <= 8191 ? 8191 - Local Pref : 0) or 0x2000 * (16 - BGP4 Import BGP4 Preference) + (Local Pref <= 0x1fff ? 0x1fff - Local Pref : 0)
RIP imported	Metric
EGP imported	Distance
Static	SR cost

8. Add up the values you have calculated.

_____

_____

_____

_____

The total is the route weight: _____

Appendix E
IP/OSPF Configuration

The IP/OSPF configuration in Figure E-1 shows an AS divided into different types of OSPF areas using several types of OSPF routers and interfaces as well as variable-length subnetting. This appendix describes parameter settings for Router 1 through Router 6. Routers W through Z and A through D are included on the network map for completion.

In the configuration in Figure E-1:

- The OSPF Area Authentication parameter is set to none for all areas.
- All Timer parameter values are left at their defaults (Hello, Dead, and Poll Interval).
- No virtual links are configured.
- Route summarization is not used.
- The Frame Relay cloud is set to the default management type. It is non-fully meshed with all Group Mode PVCs.
- There are three area border routers (Router 2, Router 3, and Router 4), two internal backbone routers (Router 1 and Router 5) and one AS boundary router (Router 6 has an Ethernet interface using RIP).
- R2 is the designated router for the FDDI segment. Router 4 must be the designated router for the Frame Relay cloud.
- The unnumbered LAN connecting Router 6 and Router D is included to ensure that every internal node in Area 0.0.0.1 is reachable from every other internal node.

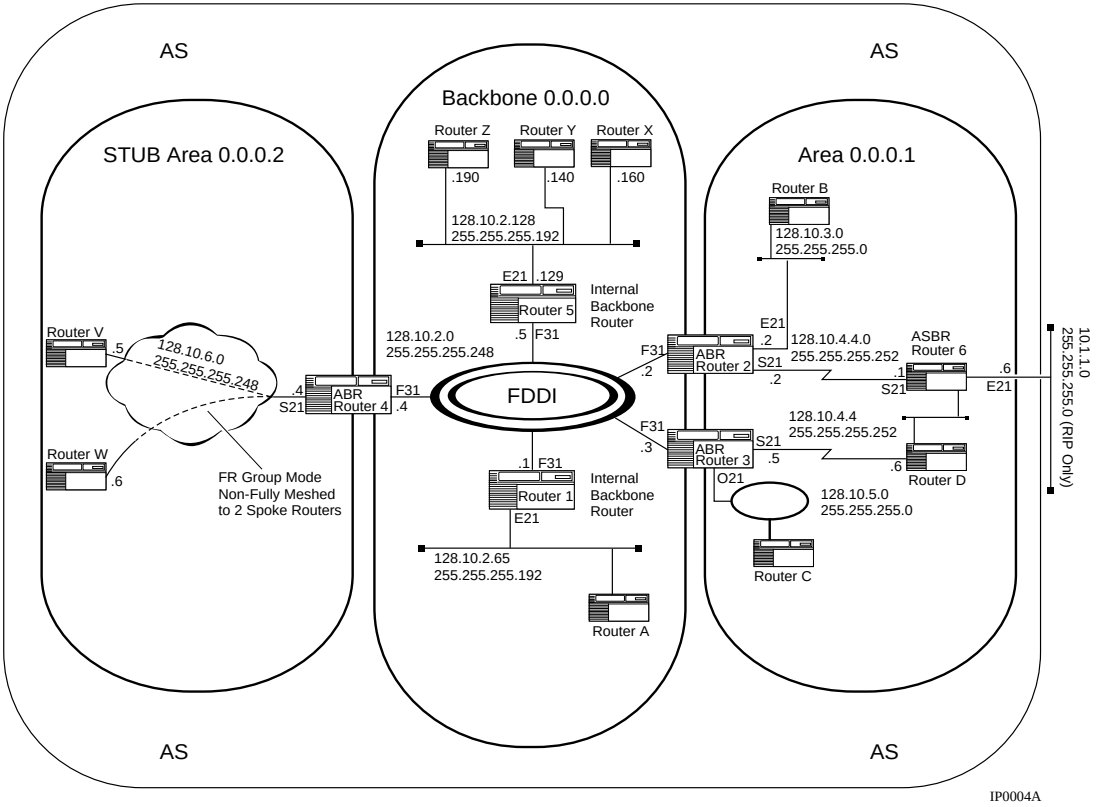

Figure E-1. IP/OSPF Configuration

The following tables list non-default configuration parameters for Router 1 through Router 6. Parameters that are not shown are set at their defaults.

Table E-1. Internal Backbone Router 1

SM Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.1
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.1
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface E21	
IP Configuration/IP Address	128.10.2.65
IP Configuration/Mask	255.255.255.192
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast

Table E-2. Area Border Router 2

SM Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.2
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.2
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
OSPF Interface/Rtr Priority	2 or greater
Interface E21	
IP Configuration/IP Address	128.10.3.2
IP Configuration/Mask	255.255.255.0
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Standard
IP Configuration/IP Address	128.10.4.2
IP Configuration/Mask	255.255.255.252
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point

Table E-3. Area Border Router 3

SM Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.3
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.3
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface O21	
IP Configuration/IP Address	128.10.5.3
IP Configuration/Mask	255.255.255.0
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Standard
IP Configuration/IP Address	128.10.4.5
IP Configuration/Mask	255.255.255.252
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point

Table E-4. Area Border Router 4

SM Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.4
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.4
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Frame Relay
IP Configuration/IP Address	128.10.6.4
IP Configuration/Mask	255.255.255.248
OSPF Area/Area	0.0.0.2
OSPF Area/AS External	No (Area 0.0.0.2 is a stub)
OSPF Interface/Broadcast Type	Point-to-multipoint standard

Note: Router 4 must be the designated router for the Frame Relay network within the cloud. To ensure this, set the Router Priority parameter on the OSPF Frame Relay interfaces for Router V and Router W to zero. The broadcast type should be set to Point-to-multipoint (standard). In addition, Router V and Router W must have IP adjacent host entries configured for each other.

Table E-5. Internal Backbone Router 5

SM Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.5
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.5
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface E21	
IP Configuration/IP Address	128.10.2.129
IP Configuration/Mask	255.255.255.192
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast

Table E-6. AS Boundary Router 6

SM Window/Parameter	Setting
Interface S21	
Wan Protocol	Standard
IP Configuration/IP Address	128.10.4.1
IP Configuration/Mask	255.255.255.252
OSPF Global/Rtr ID	128.10.4.1
OSPF Global/ AS Boundary Router	Yes
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point
Interface E21	
Add Protocols	RIP
IP Configuration/IP Address	10.1.1.6
IP Configuration/Mask	255.255.255.0

Index

A

Address Resolution Protocol
  function of, 3-27
  HP Probe, 3-29
  Inverse ARP, 3-29
  proxy ARP, 3-31
  X.25 DDN and PDN, 3-29
adjacent hosts, definition of, 3-33
aggregate route, definition of, 1-8
area, OSPF, definition of, 5-5
ARP
  see Address Resolution Protocol
ASE metric type, OSPF, 5-26
ATM (Asynchronous Transfer Mode), IP over, 3-26
autonomous system (AS)
  defintion of, 1-8

B

backbone, OSPF, definition of, 5-5
Bay Networks
  CompuServe forum, xxvi
  Customer Service FTP, xxv
  home page on World Wide Web, xxv
  InfoFACTS service, xxvii
  publications, ordering, xxii
  support programs, xxiv
  Support Source CD, xxvi
  Technical Response Center, xxiii, xxvii
  technical support, xxiii
BGP
  AS weight classes, 6-21
  AS weights, 6-21
  best route calculation, 6-7, 6-24
  interaction with OSPF, 6-25
  interior BGP (IBGP), 6-3
  Local Preference Attribute, 6-7
  message logging, 6-25
  path attributes, 6-5
  peer-to-peer communication, 6-13
BGP parameters
  BGP peers
    Connect Retry Timer, A-9
    External Advertisement Timer, A-8
    Holdtime, A-9
    Local AS to Advertise to Peer, A-10
    Peer Address, A-6
  configuration
    Local AS, A-1
  event logging
    Remote Address, A-16
  global
    BGP Identifier, A-1
    IBGP Intra AS Routing, A-3
  weight
    Weight Value, A-12, A-13, A-14
BGP-3 parameters
  BGP-3 Preference, C-16
  Enable, C-15
  Export Action, C-19
  Export Address, C-17
  Export Enable, C-19
  Export from Protocol, C-18
  Export Mask, C-17
  Export Neighbor AS, C-20
  Export Origin, C-20
  Export Peer Address, C-18

Export Peer AS, C-18
Export Use Inter AS Metric, C-19
Import Address, C-13
Import Mask, C-13
Import Peer Address, C-14
Import Peer AS, C-14
Import Peer Original AS, C-14
Import Route Origin, C-15
Preference, C-16
BGP-4 parameters
global
Enable, A-6
Blacker Front-End support, 3-26, 3-47, 3-48
addressing, 3-48, 3-49
configuring, 3-49
border router, OSPF, 5-30
broadcast address
definition of, 3-16
for subnets, 3-16
broadcast network, 5-15

C

circuitless IP interfaces, 2-15
configuring, 2-15
Classless Inter-Domain Routing (CIDR), 1-8
CompuServe, Bay Networks forum on, xxvi
configuring
circuitless IP interfaces, 2-15
Customer Service FTP, xxv
customer support. *See* getting help

D

DDN X.25 address resolution, 3-29
deleting
BGP, 2-10
BGP-3, 2-10
EGP, 2-12
IP from an interface, 2-4, 2-5, 2-6, 2-8, 2-13
dial-optimized routing for RIP, 4-9

E

EGP
modes, 7-3
overview of, 7-1
EGP parameters
Acquisition Mode, A-20
Action, C-23, C-26
Enable, C-22, C-26
Export Address, C-24
Export from Protocol, C-25
Export Mask, C-24
Export OSPF Tag, C-26
Export OSPF Type, C-25
Export Peer, C-25
Hello Timer, 7-9
Import Address, C-21
Import AS, C-22
Import Gateway, C-22
Import Peer, C-21
Interface, C-27
Metric, C-27
Preference, C-23
Enable Default Route for Subnets, A-44
equal-cost mulipath, 3-11
equal-cost multipath
OSPF, 5-13
RIP, 4-8
external route tag, OSPF, 5-28

G

getting help
from a Bay Networks Technical Response Center, xxvii
from the Support Source CD, xxvi
through CompuServe, xxvi
through Customer Service FTP, xxv
through InfoFACTS service, xxvii
through World Wide Web, xxv

H

HP Probe, definition of, 3-29

I

IBGP (interior BGP), 6-3
implementation notes
  EGP, 7-5
  OSPF, 5-7
InfoFACTS service, xxvii
interface, definition of, 3-13
Interior Gateway Protocol (IGP), 1-8
Internet Network Information Center (NIC), 1-2
Internet Requests for Comments (RFCs)
  IP router compliance, 1-17
Internet service provider (ISP) mode, 3-12
Inverse ARP, 3-29
IP
  parameters
    IP Address, A-22
    Subnet Mask, A-22
    Transmit Bcast Addr, A-23
IP accounting, 3-35
IP address
  definition of, 1-2
  network classes, 1-2
  specifying in dotted decimal notation, 1-3
IP parameters
  adjacent host
    Adjacent Host X.121 Address, A-53
    Enable, A-51
    Host Encapsulation, A-52
    IP Address, A-51
    MAC Address, A-52
    Next Hop Interface Addr, A-51
    Next Hop Interface Mask, A-52
  global
    ARP Forwarding, A-40
    Default TTL, A-41
    Estimated Hosts, A-43
    Estimated Networks, A-43
    Forwarding, A-39
    Maximum Policy Rules, A-44
    Nonlocal ARP Destination, A-41
    Nonlocal ARP Source, A-40
    RIP Diameter, A-42
    Zero Subnet Enable, A-42
interface
  Addr Mask Reply, A-26
  Address Resolution, A-27
  All Subnet Bcast, A-27
  Broadcast Address, A-25
  Checksum, A-29
  Enable, A-24
  Enable Security, A-36
  Enet Arp Encaps, A-31
  FR Broadcast DLCI, A-32
  FR Multicast DLCI#1, A-33
  FR Multicast DLCI#2, A-33
  Host Cache, A-28
  Interface Cost, A-25
  MAC Address, A-29
  Max Forwarding Table Size, A-35
  MTU Discovery, A-26
  Proxy, A-28
  Redirects, A-30
  Slot Mask, A-34
  SMDS Arp Req Address, A-32
  SMDS Group Address, A-31
  Subnet Mask, A-24
  TR Endstation, A-30
RIPSO
  Default Authority, A-62
  Default Label, A-61
  Default Level, A-62
  Enable Security, A-55
  Error Authority, A-63
  Error Label, A-63
  Implicit Authority, A-60
  Implicit Label, A-60
  Implicit Level, A-61
  Maximum Level, A-57
  May In Authority, A-59
  May Out Authority, A-58
  Minimum Level, A-57

Must In Authority, A-59
Must Out Authority, A-58
Require In Security, A-56
Require Out Security, A-56
Strip Security, A-55
static route
Address Mask, A-48
Cost, A-49
Destination IP Address, A-48
Enable, A-47
Next Hop Addr, A-49
Next Hop Mask, A-49
Preference, A-50
TFTP
Close Time Out, A-54
Default Volume, A-53
Enable, A-53
Retransmit, A-54
Retry Time Out, A-54
IP router
internal routing tables, 1-11

L

Local Preference attribute, calculating, 6-7

M

multinet
definition of, 3-15

N

NetBIOS over IP, 8-2
adding a traffic filter, 8-14
aging a cache entry, 8-9
configuring a cache, 8-7
configuring a static name, 8-12
customizing a cache search, 8-10
NetBIOS/IP parameters
global
15-Character NetBIOS Name Caching, 8-8
Enable/Disable, A-66
Rebroadcast Record Route, 8-7
interface
Enable NetBIOS Inbound Broadcasts, A-70
Enable NetBIOS Outbound Broadcasts, 8-12, A-70
Enable/Disable, A-69
NetBIOS Name Caching, A-70
Rebroadcast Address, A-71
static entry
Enable, A-71
IP Address, A-72
NetBIOS Scope ID, A-71, A-72
NetBIOS Station Name, A-72
Network Basic Input-Output System (NetBIOS) over IP, 8-2
nonbroadcast multiaccess network, 5-15

O

OSPF
area, definition of, 5-5
ASE metric type, 5-26
backbone, definition of, 5-5
backup soloist, 5-10
border router, 5-30
broadcast interface, 5-15
database synchronization, 5-3
external route tag, 5-28
features
configurable cost metrics, 5-23
link state protocol, 5-3
nonbroadcast multi-access interface, 5-15
point-to-multipoint interface, 5-16
point-to-point interface, 5-15
specifying a preferred path, 5-23
summary route, 5-34
virtual link, 5-32
OSPF parameters
area
Authentication Type, A-86
Enable, A-85
Import AS Extern, A-86
Import Summaries, A-87
Stub Metric, A-87

area range
Enable, A-89
Mask, A-89
Metric, A-90
Range Mask, A-88
Range Net, A-88
Status, A-89
export route filters
Action, C-11
Auto Tag, C-12
Enable, C-11
Export Address, C-9
Export From Protocol, C-10
Export Mask, C-10
Tag, C-12
Type, C-11
global
AS Boundary Router, A-73
ASE Metric Support, A-75
Backup Disable, A-75
Backup Log Mask, 5-14, A-76
Enable, A-73
Hold Down Timer, A-74
OSPF Slot, A-75
Primary Log Mask, A-76
Router ID, A-73
import route filters
Import Address, C-8
Preference, C-9
interface
Area ID, A-78
Dead Interval, A-81
Enable, A-78
Hello Interval, A-81
Metric Cost, A-82
MTU Size, A-83
Password, A-83
Poll Interval, A-82
Retransmit Interval, A-80
Rtr Priority, A-79
Transit Delay, A-80
Type, A-79
neighbor
Enable, A-84
Neighbor Address, A-84
Priority, A-85
virtual interface
Dead Interval, A-93
Enable, A-91
Hello Interval, A-92
Neighbors Router ID, A-90
Password, A-93
Retransmit Interval, A-92
Transit Delay, A-91

P

parameters
IP
IP Address, A-22
Subnet Mask, A-22
Transmit Bcast Addr, A-23
PDN X.25 address resolution, 3-29
point-to-multipoint network, 5-16
point-to-point network, 5-15
poisoned reverse, RIP updates, 4-3
policies, definition of, 1-13
policy parameters
Action (accept), B-4
Action (announce), B-23
Advertise (announce), B-26
Aggregator AS List (accept), B-17
Aggregator Router List (accept), B-17
Announce Tag, B-40
Apply Subnet Mask (accept), B-8, C-5
AS List (Accept), B-10
AS Path (announce), B-50
AS Path Override (announce), B-45
AS Pattern (announce), B-52
AS Weight Class (accept), B-13, B-18
Atomic (announce), B-51
BGP-3 Route Preference (accept), B-13
BGP-4 Preference (accept), B-18
Community Match (accept), B-14, B-19
Community Match (announce), B-46, B-52
EGP Interface List (announce), B-42
EGP Metric (announce), B-42

EGP Peer List (announce), B-41
Enable (accept), B-1
Enable (announce), B-20
External Route Source (announce), B-38, B-41, B-43, B-47
From BGP Peer (announce), B-35
From BGP Peer AS (announce), B-36
From EGP Peer (announce), B-32
From Gateway (accept), B-7
From OSPF Router ID (announce), B-29
From RIP Gateway (announce), B-27
Gateway List (accept), B-10
Injection List (accept), B-11, B-15
Inter-AS Metric Selector (announce), B-44
Local Preference (accept), B-18
Local Preference Override (announce), B-50
Local Preference Value (announce), B-51
Multi Exit Discriminator (announce), B-48
Multi Exit Discriminator Value (announce), B-49
Name (accept), B-2
Name (announce), B-21
Networks (accept), B-3
Networks (announce), B-22
Next Hop (announce), B-51
Origin (announce), B-45, B-49
Originating AS (accept), B-12, B-16
OSPF Metric (announce), B-40
Outbound Interface (announce), B-38
Outbound Peer AS (announce), B-47
Outbound Peer AS List (announce), B-43
Outbound Peers (announce), B-44, B-48
Peer Address (accept), B-12, B-16
Peer AS (accept), B-12, B-16
Peer List (accept), B-9
Precedence (announce), B-24
Received BGP Next Hop (announce), B-37
Received EGP Gateway (announce), B-34
Received on Interface (accept), B-7
Received on RIP Interface (announce), B-28
Received OSPF Tag (announce), B-31
Received OSPF Type (announce), B-30
Route Origin (accept), B-13, B-17
Route Preference (accept), B-5
Rule Precedence (accept), B-6
Specific Inter-AS Metric (announce), B-45
Tag (accept), B-9
Type (accept), B-8
Type (announce), B-39
preference, definition of, 1-11
Proxy ARP, 3-31

R

revised IP security option
see RIPSO
RIP
configuring timers, 4-9
poisoned reverse updates, 4-3
split horizon updates, 4-3
update mode, 4-2
RIP parameters
export route filters
Action, C-7
Enable, C-7
Export Address, C-5
Export Mask, C-6
From Protocol, C-6
Interface, C-7
Rip Metric, C-8
import route filters
Action, C-3
Enable, C-4
Import Address, C-1
Import Mask, C-2
Interface, C-3
Preference, C-4
RIP Gateway, C-2
interface
Authentication Password, A-99
Authentication Type, A-99
Broadcast Timer, A-97
Default Route Listen, A-96
Default Route Supply, A-95
Enable, A-94
Holddown Timer, A-98
Poisoned Reverse, A-96

RIP Listen, A-95
RIP Mode, A-98
RIP Supply, A-94
Time to Live, A-97
Timeout Timer, A-97
Triggered Update, A-99
RIPSO
example of, 3-43
network example, 3-44
Router Discovery
definition of, 3-46
parameters
Broadcast Type, A-64
Enable, A-64
Interface Pref, A-65
Lifetime, A-65
Maximum Interval, A-65
Minimum Interval, A-64

S

security label format, 3-39
split horizon, RIP updates, 4-3
starting
BGP, 2-9
EGP, 2-12
IP, 2-3
NetBIOS, 2-13
OSPF, 2-7
RIP, 2-5
static black hole routes
configuring, A-48, A-49
definition of, 3-46
static routes
definition of, 3-45
subnet mask
function of, 1-4
specifying, 1-4
subnets, definition of, 1-4
summary route, OSPF, 5-34
supernets, definition of, 1-7
Support Source CD, xxvi

T

TFTP (Trivial File Transfer Protocol), 3-36
timers, configuring for RIP, 4-9
Token Ring networks
ARP requests, 3-22
token ring networks
using IP over, 3-22
Trivial File Transfer Protocol
function of, 3-36

U

unnumbered interface
definition of, 2-16
update mode, RIP, 4-2

V

virtual link, OSPF, 5-32

W

weight, definition of, 1-11
World Wide Web, Bay Networks home page on, xxv